



Experiences with Tools at NERSC

Richard Gerber
NERSC User Services

Programming weather, climate, and earth-system models
on heterogeneous multi-core platforms

September 7, 2011 at the National Center for Atmospheric Research in Boulder, Colorado



U.S. DEPARTMENT OF
ENERGY

Office of
Science



National Energy Research
Scientific Computing Center



Lawrence Berkeley
National Laboratory



- Thanks for the invitation
- My professional goal is to enable scientists to use HPC easily and effectively
- Contribute to important discoveries about how our natural world works
- Make a difference
- So it is an honor & meaningful to me to participate in this conference
- One of my primary roles is as deputy on our next procurement team & we are extremely interested in learning about your experiences with hybrid systems and programming



Outline

- Recent experiences providing tools
- Observations about tool usage by NERSC users
- What I'd like to see in tools for the HPC community



Focus

- I'll focus on performance tools
 - Who is the audience?
 - Scientists?
 - CS code engineers?
 - Scientist/engineers?
- Development tools
 - Eclipse: We considered installing this, but concluded users best served by installing themselves in their own directories
 - Compilers: PGI, Cray, HMPP directives-based (your experiences?)
- Debuggers
 - Allinea DDT & Rouge Wave Totalview
 - Both are here and will present today
 - Both are installed and supported at NERSC
 - Generally good experiences, but not extensively used by users



Disclaimers

- These are my observations and experiences
- Not a review or critique of any given tool
- Not comprehensive survey of tools



Genesis of This Talk

Hi Richard:

I started looking at the single-core performance on Hopper. The first problem that I encountered was that running the code with <tool name redacted> causes it to crash; so I gave up on that approach. What other profiling tools are you planning to support? I did look at gprof on the code but that also seems to have problems.

Chris



The Plan

- Provide Chris with tools to scale and test his code's (HiRAM?) performance on NERSC's Cray XE6 Hopper (#5 in Top 500)
- At the same time test and install some tools to benefit all NERSC users
- Report on my successes and survey the tools we would have available at NERSC (including our test GPU system)



Experience with Chris Kerr

- Chris & I went back and forth over a number of months, trying to install and use various tools.
- All failed
 - Couldn't build a tool with some specific compiler
 - Tool crashed (we did find some bugs)
 - Tool appeared to run, but produced no output
- Chris found a solution for his problem, but it was a lot of work
 - Gettimeofday(), print, somebody's private tool
- My result: no success identifying and installing new tools for our entire user base



Experience with NERSC Users

- NERSC has about 4,000 users
 - All levels of sophistication and experience
 - We're committed to supporting both the cutting edge & production HPC computing for the masses
- Users often ask for advice on which tools to use and we give them suggestions
- Our experience is that very few use programming/debugging/development tools
- A few users use a few tools a lot, but many try a tool only once

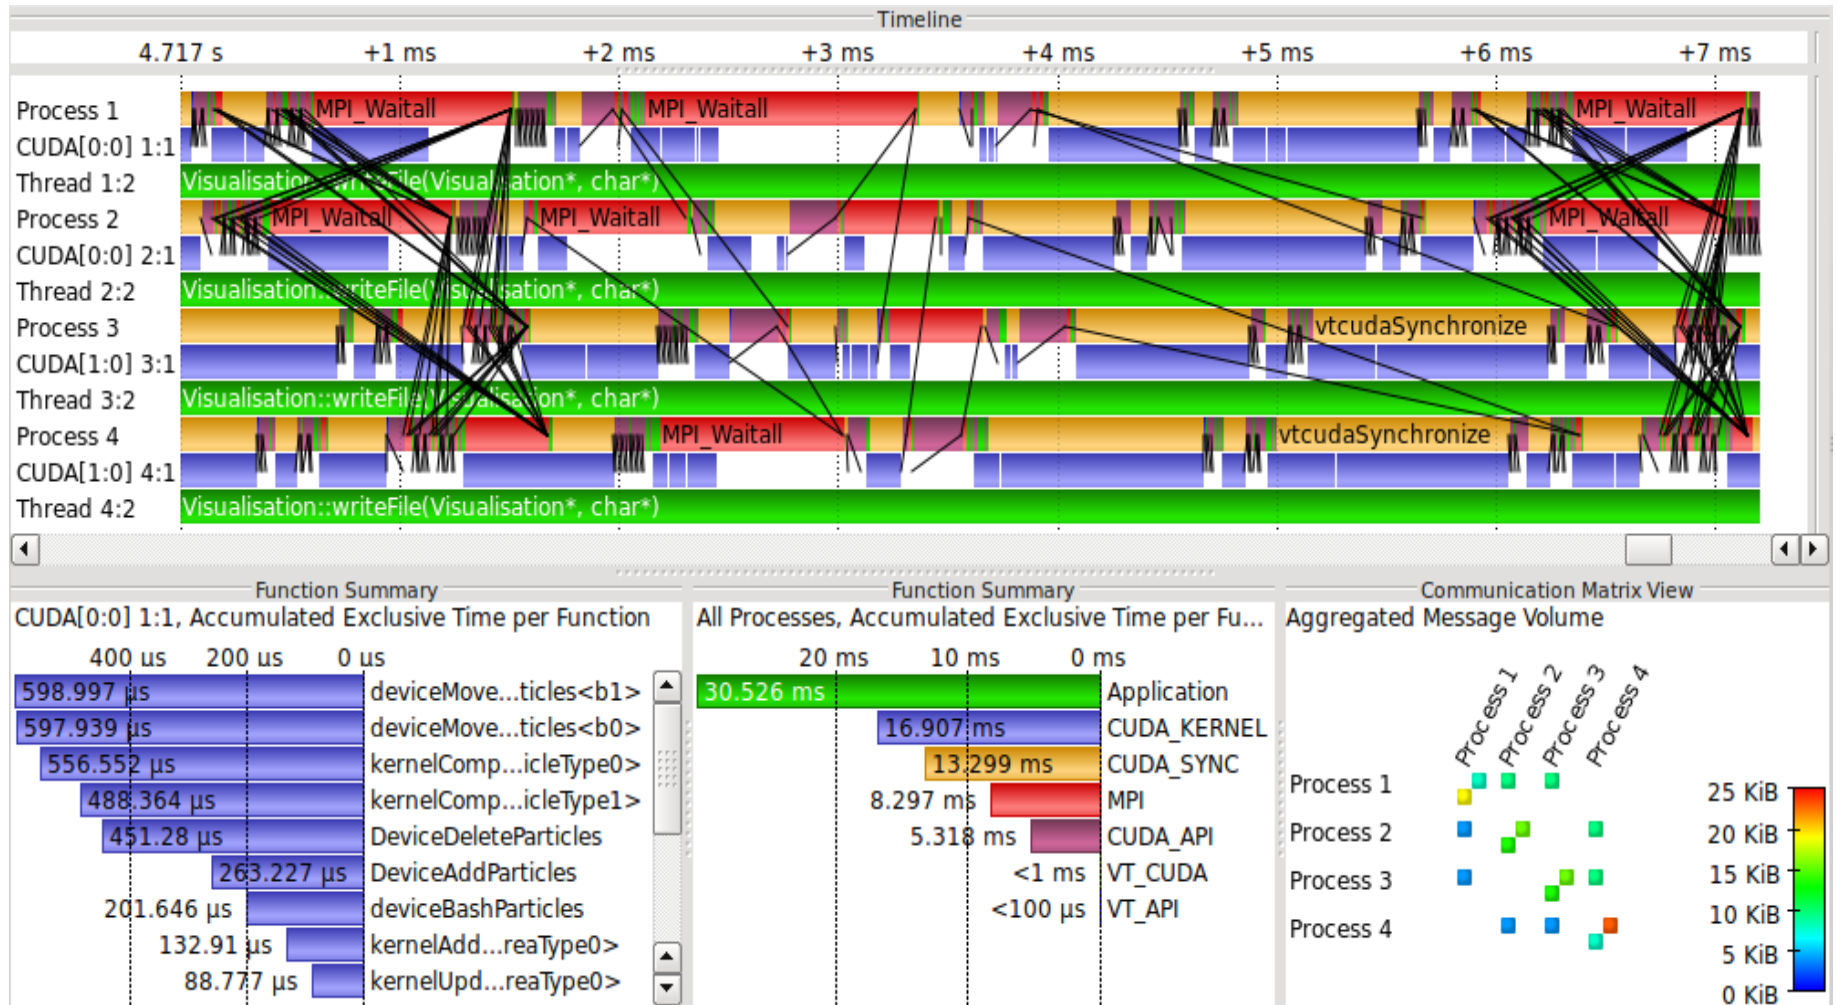


Why?

- Extremely effective?
- More likely: Too confusing, difficult, didn't work, don't know how to use, don't know which to use
- It's not that we don't have tools that address specific issues
 - GPU/CUDA tools & compilers
 - TAU, PAPI, HPC Toolkit
 - Craypat, IBM HPC tools, OpenSpeedShop
 - Valgrind
 - Vampirtrace
- But do most users have the resources to learn how to use these tools, esp. when they don't know if there will be any benefit from any given one?



Vampir w/ CUDA





Ease of Use

Use

Follow these **10 STEPS** to perform the basic analysis of your program using a performance analysis tool, not a debugging tool, start with a fully debug capable of running to a planned completion or an intentional termination environment modules first. This ensures that the correct links and libraries



Users Want (Need?) Tools

- Users are asking for tools because HPC systems and programming models are changing
- More and more components to worry about
 - CPU (caches, FPUs, pipelining, ...)
 - Data movement to main memory, GPU memory, levels of cache
 - I/O
 - Network (message passing)
 - CPU Threads (OpenMP)
 - GPU performance



Questions to You

- What tools do you use?
- What tools do you want?
- What would you like centers to support?
- Can you get to exascale without tools?



What I Want in a Tool

- Let the users help themselves
- Work for everyone all (most of?) the time
- Easy to use
- Useful
- Easy to interpret the results
- Affordable (\$\$ or manpower support costs)
- Simple, supplement existing complex tools
 - Point the way for a “deeper dive” in problem areas



IPM as A Model

- NERSC's IPM: Integrated Performance Monitoring
- How it works (user perspective)
 - % module load IPM*
 - Run program as normal
 - Look at results on the web
- It's that easy!
 - And extremely low overhead, so IPM is examining your production code

* (As long as your system supports dynamic load libs)



What IPM measures

- IPM “only” gives a high-level, entire-program-centric view
- Still, very valuable guidance
 - Shows whole-run info per MPI task, OpenMP thread, (CUDA under development)
 - Many pieces of data in one place
- Reveals what many users don’t know about their code
 - High-water memory usage (per task)
 - Load balance
 - Call imbalance
 - MPI time
 - I/O time



IPM w/ CUDA: Accuracy

Benchmark	Kernel Invocations	GPU Kernel Execution Time (sec)		
		CUDA Profiler	IPM	Difference (%)
BlackScholes	512	2.540677	2.543700	0.12
FDTD3d	5	0.101354	0.101550	0.19
MersenneTwister	202	1.126475	1.127000	0.05
MonteCarlo	2	0.001988	0.002025	1.87
concurrentKernels	9	0.613755	0.614000	0.04
eigenvalues	300	5.328266	5.331000	0.05
quasirandomGenerator	42	0.039536	0.039736	0.51
scan	3300	1.412912	1.430200	1.22

- Benchmarks from the CUDA SDK
 - Run on 1 node of Dirac @ NERSC
 - 2x Nehalem quad core, 1x NVIDIA Tesla C2050 (“Fermi”) GPU
 - Comparison of the results from the CUDA profiler with IPM
 - Very good agreement of the results
 - IPM timings always larger than CUDA profiler (bracketing)



IPM Examples

Click on your job from a list on the NERSC web site.

NERSC Completed Batch Jobs Listing

1-800-66-NERSC, option 3 or 510 486-8611

Off-Hours Status & Passwords
1-800-66-NERSC, option 1

Send us feedback about this page

Displaying 302 jobs that completed between May-30-11 00:00 and Sep-05-11 23:59

Show All entries Search:

#	Host	JobID	Job Name	User	Nds	Complete	Wall hrs	Raw hrs
166	hopper	550228	hs65536	colliera	2,731	06/27/11 20:08	0.160	10,468.83
241	hopper	485980	18Ne_7	shirokov	2,097	06/13/11 23:03	0.404	20,312.94
220	hopper	487726	18Ne_7	shirokov	2,097	06/14/11 16:53	1.460	73,464.90
175	hopper	544457	hs32768	colliera	1,366	06/25/11 09:12	0.105	3,433.21
176	hopper	544458	hs16384	colliera	683	06/25/11 08:38	0.076	1,247.61
192	hopper	543656	GTS-6-thre... [Full Name]	pravn	512	06/24/11 16:55	0.171	2,106.03
130	hopper	601346	MFDn	pmaris	504	07/08/11 15:09	0.477	5,765.76
299	hopper	442128	ryd_ion3d	turker	500	05/30/11 12:45	0.419	5,026.67
298	hopper	442211	ryd_ion3d	turker	500	05/30/11 14:02	0.431	5,176.67
297	hopper	442271	ryd_ion3d	turker	500	05/30/11 15:11	0.456	5,466.67
296	hopper	442593	ryd_ion3d	turker	500	05/30/11 15:59	0.427	5,123.33
295	hopper	442955	ryd_ion3d	turker	500	05/30/11 16:56	0.420	5,036.67
294	hopper	442991	ryd_ion3d	turker	500	05/30/11 16:56	0.418	5,020.00

View Page in: [CMS](#) [Draft Site](#) [Published Site](#) Logged in as Richard Gerber - [Log out](#)



U.S. DEPARTMENT OF ENERGY

Office of Science



Lawrence Berkeley National Laboratory



IPM Examples

Click on the metric you are want.

NERSC job details
http://www.nersc.gov/REST/jobs/job_details.php?stepid=732423.sdb×tamp=1313679078&completion=1313679081

IPM Summary

Executable					./wrf.exe
Number of tasks	512	Aggregate GFlop/sec	0.1482	Average GFlop/sec/task	0.0003
Average wall secs	8.861e+01	Aggregate memory (GB)	32.2626	Average memory/task (GB)	0.0630
Average MPI secs/task	7.898e+01	MPI time %	89.14	Aggregate MPI calls made	7.027e+07

IPM Summary Statistics - 512 tasks

Metric	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum
Aggregate Floating Point Operations (Flop x 10**9)	1.313e+01	2.565e-02	6.10	1.713e-02	2.758e-02
GFlop/sec	1.482e-01	2.895e-04	6.10	1.934e-04	3.114e-04
Maximum Memory Usage (GBytes)	3.226e+01	6.301e-02	10.12	5.701e-02	1.947e-01
Time Spent in MPI Routines (sec)	4.044e+04	7.898e+01	4.05	9.801e+00	8.359e+01
Wallclock Time (sec)	4.537e+04	8.861e+01	0.10	8.848e+01	8.895e+01

Memory in units of gigabytes; time in seconds.

Hardware counter statistics - 512 tasks

Counter Name	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum
PAPI FP OPS	1.161799e+12	2.269139e+09	6.09	1.515023e+09	2.439529e+09

MPI Time Statistics - 512 tasks

Call	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum	% of MPI	% of wall
MPI Bcast	3.517e+04	6.869e+01	4.48	4.342e-01	7.269e+01	86.969	77.520
MPI Scatterv	2.589e+03	5.057e+00	5.79	1.059e+00	5.540e+00	6.403	5.707
MPI Wait	2.176e+03	4.249e+00	17.82	1.250e+00	4.968e+00	5.380	4.795
MPI Gatherv	4.312e+02	8.422e-01	36.44	3.552e-03	2.271e+00	1.066	0.950
MPI Isend	5.250e+01	1.025e-01	11.96	7.182e-02	1.259e-01	0.130	0.116
MPI Irecv	1.033e+01	2.017e-02	10.21	1.217e-02	2.613e-02	0.026	0.023
MPI Gather	1.021e+01	1.995e-02	502.07	1.391e-03	1.434e+00	0.025	0.023
MPI Comm_rank	4.563e-01	8.913e-04	4.74	7.799e-04	1.404e-03	0.001	0.001
MPI Comm_size	9.629e-02	1.881e-04	10.65	1.462e-04	4.859e-04	0.000	0.000
MPI Init	0.000e+00	0.000e+00	0.000e+00	0.000e+00	0.000	0.000	
MPI Finalize	0.000e+00	0.000e+00	0.000e+00	0.000e+00	0.000	0.000	

Average MPI Time per Task



- MPI_Bcast
- MPI_Scatterv
- MPI_Wait
- MPI_Gatherv



U.S. DEPARTMENT OF
ENERGY

Office of
Science



Lawrence Berkeley
National Laboratory



IPM Examples

NERSC job details

http://www.nersc.gov/REST/jobs/ipm_summary.php?stepid=627129.sdb&name=memory×tamp=1311046935

Task distribution of *Maximum Memory Usage (GBytes)* - as a percentage of maximum

The MPI rank is the sum of the column and row indices in the table.

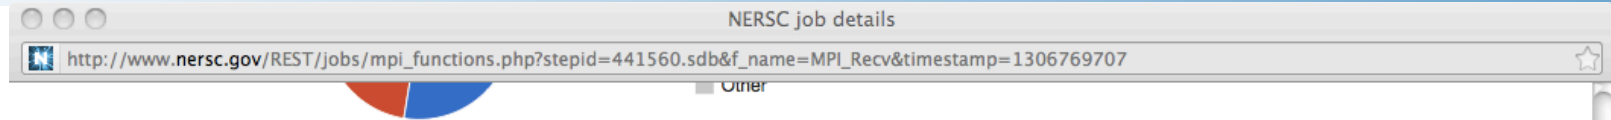
Table Columns: 32

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	67	68	67	69	67	69	68	70	68	69	69	70	69	70	71	71	70	72	69	71	71	73	70	72	71	72	71	72	72	74	72	73	
32	69	71	70	72	70	72	70	72	71	72	72	73	71	72	72	73	73	75	72	74	73	75	72	74	75	76	74	75	74	75	74	75	
64	71	71	72	73	72	73	73	74	72	72	73	73	73	74	74	74	75	74	75	76	77	75	76	76	76	75	75	77	77	77	76		
96	74	75	75	76	74	75	75	76	75	76	76	76	75	75	76	75	78	79	77	79	78	79	77	78	79	79	78	78	79	78	78		
128	86	88	86	87	87	89	86	89	88	89	87	88	88	89	88	89	86	87	85	87	87	89	86	88	87	88	86	87	88	89	88	89	
160	89	91	89	91	89	91	88	90	91	92	90	91	90	91	90	91	89	90	88	90	88	90	88	90	90	91	90	91	90	91	89	90	
192	95	96	95	96	96	97	96	97	96	96	95	95	97	97	97	96	95	95	94	95	96	97	95	96	95	95	95	95	97	97	96	96	
224	98	99	98	98	98	99	97	98	100	99	98	98	99	99	98	97	98	99	97	98	98	99	97	98	99	99	99	98	97	99	98	98	97
256	70	72	71	73	70	72	71	73	72	73	72	73	71	72	73	73	73	75	73	74	73	75	72	74	75	76	74	75	74	76	74	75	
288	67	69	68	70	66	68	66	68	69	70	71	67	68	68	69	71	73	70	72	69	71	68	70	72	73	72	73	70	71	70	71		
320	74	75	76	76	74	75	75	77	75	75	76	76	75	75	76	78	79	77	78	78	79	77	79	80	79	79	78	79	79	78	78		
352	72	73	73	74	70	71	71	72	72	72	73	73	70	70	72	71	75	76	75	76	74	74	73	74	76	76	76	75	74	74	73	73	
384	90	92	89	91	90	91	89	91	91	93	91	92	91	92	91	91	89	91	89	90	89	91	88	90	91	92	90	91	90	91	90	91	
416	87	89	87	89	85	87	85	86	89	90	88	89	87	88	86	87	86	88	86	88	85	87	84	86	88	89	87	88	86	87	86	87	
448	99	100	98	99	99	100	98	99	100	100	99	99	100	100	99	99	99	99	98	99	98	99	98	99	99	99	99	99	99	99	99	98	
480	96	97	95	96	94	95	93	94	97	97	96	96	95	95	94	94	96	97	95	96	94	95	93	94	97	96	96	95	95	94	94	94	
512	71	72	72	73	72	73	73	74	71	72	71	72	72	73	73	74	75	75	74	75	76	77	75	76	74	75	73	74	75	76	75	76	
544	74	75	75	76	74	75	75	76	74	75	75	76	74	75	75	76	78	79	77	78	78	79	77	78	77	78	77	78	77	78	77	78	
576	68	69	70	70	69	70	70	72	68	67	69	68	69	68	71	69	72	73	71	72	73	74	72	74	72	70	71	70	73	72	73	71	
608	72	73	73	74	72	72	72	74	71	70	72	71	71	70	72	71	76	77	75	76	76	76	74	76	75	74	74	73	75	74	74	73	
640	94	95	94	94	95	96	94	95	94	94	93	93	94	94	94	94	94	93	94	95	96	94	95	93	94	92	93	94	95	93	95		
672	98	98	97	97	97	98	96	97	97	98	96	97	96	97	96	97	97	98	96	97	97	98	96	97	96	97	96	97	96	97	96	97	
704	88	89	87	88	89	90	89	90	88	86	86	86	89	88	88	87	88	89	87	88	89	90	89	90	88	86	87	86	89	88	88	87	
736	91	92	91	91	91	92	90	91	91	90	90	89	91	89	90	88	91	92	91	92	91	92	90	91	91	90	90	89	91	89	90	88	
768	75	76	76	77	75	75	76	76	74	75	75	76	74	75	75	76	78	79	78	78	79	77	78	78	79	77	78	77	78	77	78		
800	72	73	73	74	71	71	71	72	71	73	72	74	70	70	71	72	76	76	75	76	73	74	73	73	75	76	74	75	73	73	72	73	
832	72	72	73	74	71	72	72	74	71	70	72	71	71	70	72	71	75	76	75	76	75	76	75	76	75	74	75	73	75	74	74	73	
864	69	70	70	71	67	68	68	69	69	67	69	68	66	65	68	66	73	74	72	73	71	72	70	71	72	71	71	71	70	69	69	68	
896	99	99	98	98	98	99	97	98	98	99	97	97	97	98	96	97	98	99	97	98	98	99	97	98	97	98	96	97	96	98	96	97	
928	95	96	95	95	93	94	92	93	95	96	94	95	92	93	92	93	95	96	94	95	93	94	92	93	94	95	94	95	92	93	92	93	
960	92	93	91	92	91	92	91	92	92	90	91	90	92	90	91	89	92	93	91	92	91	92	91	92	92	90	91	90	91	90	91	89	
992	89	90	89	89	87	88	87	88	89	87	88	87	87	86	86	85	89	90	88	90	88	88	87	88	89	87	88	87	87	85	86	84	
1024	87	88	86	87	87	90	87	89	88	89	87	88	89	90	88	89	86	88	85	87	87	89	86	89	87	88	86	87	88	90	88	89	
1056	90	92	89	91	89	91	89	90	91	92	91	92	91	92	90	91	89	91	88	90	89	91	88	90	91	92	90	91	90	91	90	91	
1088	95	96	95	95	96	97	96	97	96	96	95	95	97	97	97	96	95	96	94	95	96	97	96	96	95	96	95	95	97	97	97	96	
1120	98	99	98	99	98	99	97	98	100	99	99	99	99	99	98	98	98	99	97	98	98	99	97	98	99	99	98	98	99	98	98	97	
1152	71	73	71	73	73	74	72	74	73	74	72	74	74	75	74	74	69	71	69	70	70	72	69	71	71	72	70	71	72	73	71	72	
1184	75	77	74	76	75	76	74	76	76	77	75	77	76	77	76	76	72	74	71	73	72	74	71	73	74	75	73	74	74	75	73	74	
1216	76	77	76	76	78	79	76	78	77	77	76	76	78	78	77	74	75	73	74	75	76	75	75	75	75	74	73	76	76	75	74		
1248	80	80	79	80	79	81	79	80	81	81	80	80	81	81	79	80	77	78	77	78	77	78	77	78	79	77	77	78	78	77	76		
1280	90	92	90	92	90	92	89	91	92	93	91	92	91	92	91	89	92	89	91	89	92	89	91	91	92	90	92	91	92	90	91		
1312	87	89	87	89	86	87	85	86	89	90	88	90	87	88	86	87	87	89	86	88	85	87	85	86	88	89	87	89	87	88	86	87	
1344	99	100	98	100	99	100	98	99	100	100	99	99	100	100	99	99	99	99	98	99	98	99	98	99	98	99	100	100	99	99	99	99	
1376	96	97	95	96	94	95	93	94	97	97	96	96	95	95	94	94	96	97	95	96	94	95	93	94	97	97	96	96	95	95	94	94	
1408	75	77	74	77	75	77	74	76	77	78	76	77	76	77	76	77	73	75	72	74	73	74	72	74	75	76	74	75	74	74	74		
1440	72	75	72	74	71	72	70	72	74	75	74	74	72	73	72	72	70	72	69	71	68	70	67	70	72	72	71	72	70	71	69	70	
1472	80	81	80	80	80	81	79	80	81	81	79	80	80	80	80	79	78	79	77	78	77	79	77	78	79	77	77	78	78	77	77		
1504	77	78	77	78	75	76	75	76	79	78	77	77	77	76	76	75	75	76	75	75	73	74	72	73	76	75	75	74	74	73	72		
1536	94	95	94	94	95	96	94	95	94	94	93	94	95	96	94	95	94	95	93	94	95	96	94	95	93	94	92	93	94	95	94	95	





IPM Examples



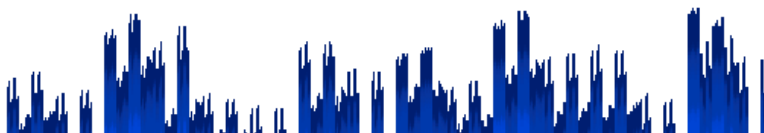
Time spent by each task in *MPI_Recv* as a percentage of the maximum value

The MPI rank represented by each cell in the table is the sum of the cell's column and row indices.

Table Columns: 32

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	74	76	69	70	77	77	70	71	60	62	60	61	64	65	65	64	78	79	72	72	78	79	73	73	63	64	63	64	66	65	66	66
32	70	71	63	65	72	72	65	66	52	54	53	55	56	57	58	58	71	73	67	68	72	73	67	69	55	55	54	55	58	58	57	57
64	91	92	88	90	91	93	90	91	76	77	74	76	79	81	79	79	95	98	92	96	98	98	96	96	78	81	77	80	84	84	83	82
96	86	87	80	80	86	89	81	82	62	63	61	61	65	67	65	65	91	94	83	86	94	94	87	86	64	66	63	65	70	69	69	68
128	69	71	64	65	70	72	65	66	54	56	53	55	60	60	59	59	69	70	64	65	69	70	65	66	56	57	55	56	60	59	58	58
160	65	67	57	59	67	68	60	61	42	43	43	44	47	46	47	47	66	67	58	60	67	67	60	60	43	44	45	46	46	46	48	48
192	86	87	80	81	88	89	82	83	67	68	66	67	71	72	70	70	85	88	81	84	89	88	84	84	68	71	66	69	74	73	72	71
224	79	79	71	71	80	80	72	72	49	49	49	50	52	52	53	52	76	79	70	73	79	79	73	74	50	52	48	52	55	55	54	54
256	83	84	81	83	85	85	84	85	69	71	69	70	74	75	74	74	85	86	85	87	86	87	86	87	73	73	71	71	76	75	74	73
288	72	73	66	68	73	74	69	70	60	62	59	60	64	65	63	63	75	76	69	71	76	76	71	71	63	63	61	61	65	65	64	64
320	94	95	93	94	93	96	94	95	77	78	75	75	80	81	78	78	99	99	96	96	99	99	98	97	79	80	77	78	83	82	81	80
352	82	83	74	75	83	84	75	76	61	62	66	66	65	65	69	69	84	85	76	77	85	85	77	78	63	64	66	67	66	66	69	69
384	84	86	76	78	86	88	79	80	64	64	63	64	68	68	67	67	85	86	77	78	85	86	78	79	65	65	66	66	68	67	69	68
416	69	70	60	62	71	72	63	64	45	46	46	47	49	50	50	51	68	69	61	62	69	70	62	62	47	47	46	47	49	49	47	49
448	98	98	99	100	99	98	100	100	85	85	78	77	89	89	81	80	94	94	95	96	94	94	97	97	87	87	78	79	91	90	82	81
480	82	82	72	73	82	84	74	74	54	54	52	52	56	57	54	55	83	83	72	73	84	84	71	73	54	55	53	54	57	57	55	56
512	88	89	89	91	90	91	92	93	78	80	78	80	83	84	84	84	93	94	93	95	94	94	95	94	80	81	82	83	85	84	86	85
544	90	91	79	81	91	93	82	83	63	65	63	64	67	69	68	68	96	97	84	84	98	97	85	85	66	66	66	67	69	68	70	69
576	76	77	74	75	76	77	76	76	65	65	66	67	67	67	70	70	79	82	77	80	83	82	81	81	67	69	69	71	72	72	74	73
608	76	77	69	69	76	78	70	71	57	57	57	58	60	59	60	60	79	81	72	75	82	81	75	75	60	62	58	61	64	63	63	62
640	92	94	85	87	96	96	90	89	75	77	75	77	81	81	81	81	93	94	87	88	94	94	88	88	77	78	76	77	81	80	80	79
672	86	87	78	78	88	88	80	80	57	58	58	59	62	62	63	63	87	87	78	79	88	87	79	79	59	59	60	60	62	61	63	63
704	78	79	74	75	80	80	76	75	63	64	63	63	68	67	67	66	78	81	74	77	81	81	77	77	64	67	63	66	70	69	69	67
736	70	71	65	65	71	71	66	67	52	52	50	51	55	54	53	53	71	73	64	67	74	73	68	68	53	55	52	54	57	57	56	56
768	94	96	87	88	96	98	90	90	71	73	70	71	75	77	75	75	95	95	87	87	96	96	87	88	74	75	73	74	79	78	77	77
800	79	80	70	71	81	83	72	73	57	58	62	63	62	62	67	67	80	79	71	71	80	80	71	72	60	60	64	64	63	63	66	67
832	82	83	77	77	82	84	78	78	63	64	64	64	66	66	67	66	86	86	79	80	87	86	81	80	65	66	64	65	69	68	68	67
864	69	70	65	65	70	70	65	66	56	56	56	56	59	59	58	58	73	73	67	67	73	73	67	67	58	59	58	59	61	59	61	59
896	92	92	85	86	93	94	88	88	71	72	69	70	76	75	74	74	92	92	87	88	94	93	88	89	73	73	70	71	76	76	74	74
928	73	74	64	65	75	76	66	67	43	44	48	49	47	48	53	54	74	73	64	65	73	74	64	66	45	45	51	52	47	48	53	54
960	74	74	70	70	75	75	71	71	63	64	61	61	66	66	64	63	76	77	72	73	77	77	73	73	63	64	62	62	66	66	66	64
992	63	63	57	57	63	64	58	58	41	41	42	42	43	42	44	43	61	61	56	57	62	62	56	57	41	41	42	42	43	42	43	43

Time vs. MPI Rank for *MPI_Recv*





Summary

- HPC tools are, in general, difficult to use or require a large investment of time to learn to use and interpret the results
- There is a need for tools to help users get to exascale (or even petascale).
- Simple, easy to use tools would be extremely useful, even if they did not give low-level details