


Performance Update

E3SM 2019 March Madness Edition



Phil Jones, Sarat Sreepathi on behalf of Performance Group:

Oksana Guba, Mathias Jacquelin, Noel Keen, Jayesh Krishna, Azamat Mametjanov,
Henry Moncada, Hongzhang Shan, Mark Taylor, Pat Worley, Min Xu

Goals

- Maximize throughput for science sims

Goals

- Maximize throughput for science sims
- Prepare for new architectures
 - CompyMcNodeFace (Skylake)

Goals

- Maximize throughput for science sims
- Prepare for new architectures
 - CompyMcNodeFace (Skylake)
 - Summit (P9/GPU)
 - Perlmutter (Better Call it Saul, EPYC/GPU, Shasta/slingshot, no spinning disk)
 - A21 (newly announced Intel Xeon SP with Xe GPU, Shasta/slingshot, memory innov)
 - Trend?



Goals

- Maximize throughput for science sims
- Prepare for new architectures
 - CompyMcNodeFace (Skylake)
 - Summit (P9/GPU)
 - Perlmutter (Better Call it Saul, EPYC/GPU, Shasta/slingshot, no spinning disk)
 - A21 (newly announced Intel Xeon SP with Xe GPU, Shasta/slingshot, memory innov)
 - Trend?
- Monitor/measure performance
 - Standard benchmarks
 - Performance tools





Performance Analytics for Climate Experiments

Infrastructure to provide executive summary of experiments performance.



- Harvest/collect performance of all sims
 - Also collateral data to sell
- Central hub of performance data
- Interactively deep-dive as desired
- Facilitate performance research
- User input, existing scripts



<https://pace.ornl.gov>



Summary Performance Graphs



Atmosphere model time distribution

Load Balance: MPI Task Mapping

Status

- Phase 2 Baseline Benchmarks
 - E3SM metric to track over time
 - How best to present/utilize this data?
- Ongoing work high-res
 - MPAS, RRTMG
 - Threading (MPAS)
- Initial profiles/optimization of BGC configuration
 - I/O bound
- I/O
 - PIO2, ADIOS
 - Substantially better performance
- KNL dead to us
 - Well, maybe some vector/thread improvements still needed

Cori-KNL

nodes	cores	Best SYPD ATM
85	68 per node	0.41
169		0.79
338		1.5
675		2.66
1350		3.94
2700		5.2

nodes	git hash: 1b2b515 10/2017	Master 12/2018	Maint-1.0 12/2018
85	0.41	0.40	0.35
169	0.79	0.75	0.65
338	1.50	1.42	1.27
675	2.66	2.46	2.14
1350	3.94	3.71	3.08

Latest update: 2019-02-17 (Collected by @Patrick Worley). Code: v1 maintenance

nodes	cores	OCN SYPD	ICE SYPD	CPL SYPD	hyperthreading
64	4096	0.59	1.53	2.37	yes (2-way, 4096x2)
75	4800	0.67	1.73	1.60	yes (2-way, 4800x2)
100	6400	0.82	1.95	2.38	yes (2-way, 6400x2)
128	8192	1.02	2.54	3.20	yes (2-way, 8192x2)
150	9600	1.10	2.80	1.30	yes (2-way, 9600x2)
200	12800	1.30	3.27	2.40	yes (2-way, 12800x2)
300	19200	1.67	3.90	1.21	yes (2-way, 19200x2)
400	25600	1.26	4.35	1.70	no (25600x1)
500	32000	2.21	4.81	1.57	no (32000x1)
750	48000	2.61	4.98	1.72	no (48000x1)
1000	64000	3.13	4.53	0.71	no (64000x1)

Summit – no GPU

Summit

High-res F case without I/O

Latest update: 2019-02-04. Code: v1 maintenance branch. Compiler: PGI/18.10

Nodes	Cores	Threads	smt	SYPD ATM	SYPD LND
100	8400	1	2	0.37	183.41
200	16800	1	2	0.78	256.02
400	33600	1	2	1.38	473.05
515	43200	1	2	1.74	407.70

Summit


High-res G case without I/O

Latest update: 2019-02-04. Code: v1 maintenance branch. Compiler: PGI/18.10

nodes	cores (cpl,ice,ocn)	cores (others)	threads	smt	SYPD OCN	SYPD ICE	SYPD CPL
115	9600	9600	1	2	1.41	4.17	6.91
191	16000	16000	1	2	2.01	6.22	9.42
572	48000	16000	1	2	4.14	10.40	2.48

GPU Strategy

- Short term: 737-MAX model
 - Not a different model
 - No need for additional training
 - Porting exercise
 - Software to hide details
 - Kokkos
 - OpenACC/OpenMP
 - Issues
 - How much to port (eg v1 physics)
- Longer term
 - Utilize GPU differently
 - Subgrid models (SAM, ocean ML)
 - Task parallelism to split work across devices



Related sessions

- Performance breakouts (3p today, 10a tomorrow)
 - Input on how to best use/present perf data
 - Input on v1 physics
- Discussions in SCREAM, Algorithm NGD
- Remember Deep Dives 4:30 tomorrow