



# Star-CCM+

## Performance Benchmarking and Profiling

Nov 2018

- **The following research was performed under the HPC Advisory Council activities**
  - Compute resource - HPC Advisory Council Cluster Center
- **The following was done to provide best practices**
  - Star-CCM+ performance overview over AMD EPYC based platforms
  - Understanding Star-CCM+ communication patterns
- **More info on Star-CCM+**
  - <https://mdx.plm.automation.siemens.com/star-ccm-plus>


- **Computational Fluid Dynamics (CFD)**
  - Enables the study of the dynamics of things that flow
  - Enable better understanding of qualitative and quantitative physical phenomena in the flow which is used to improve engineering design
- **CFD brings together a number of different disciplines**
  - Fluid dynamics, mathematical theory of partial differential systems, computational geometry, numerical analysis, Computer science
- **Star-CCM+ is a leading CFD application from Siemens**
  - Widely used in almost every industry sector and manufactured product.

## Venus cluster


- **Supermicro AS -2023US-TR4 8-node cluster**
  - Dual Socket AMD EPYC 7551 32-Core Processor @ 2.00GHz
  - Mellanox ConnectX-5 EDR 100Gb/s InfiniBand
  - Mellanox Switch-IB 2 SB7800 36-Port 100Gb/s EDR InfiniBand switch
  - Memory: 256GB DDR4 2677MHz RDIMMs per node
  - 240GB 7.2K RPM SSD 2.5" hard drive per node
- **Software**
  - OS: RHEL 7.5, MLNX\_OFED 4.4
  - MPI: Intel MPI
  - Star-CCM+ : 13.04.010



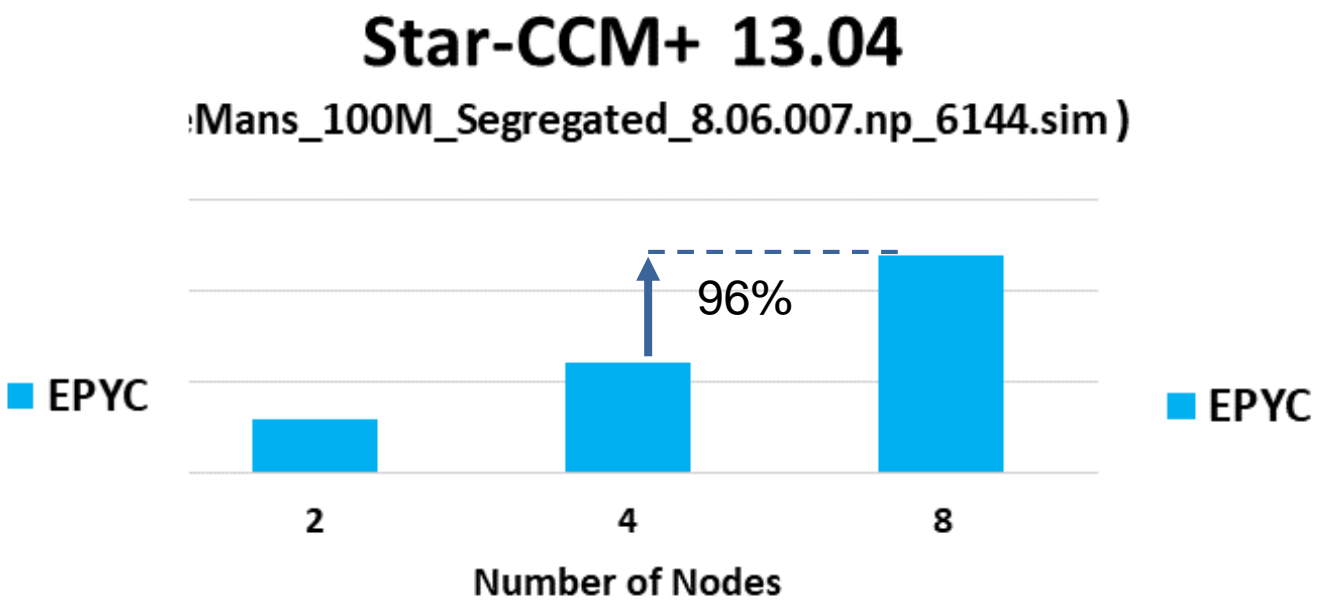
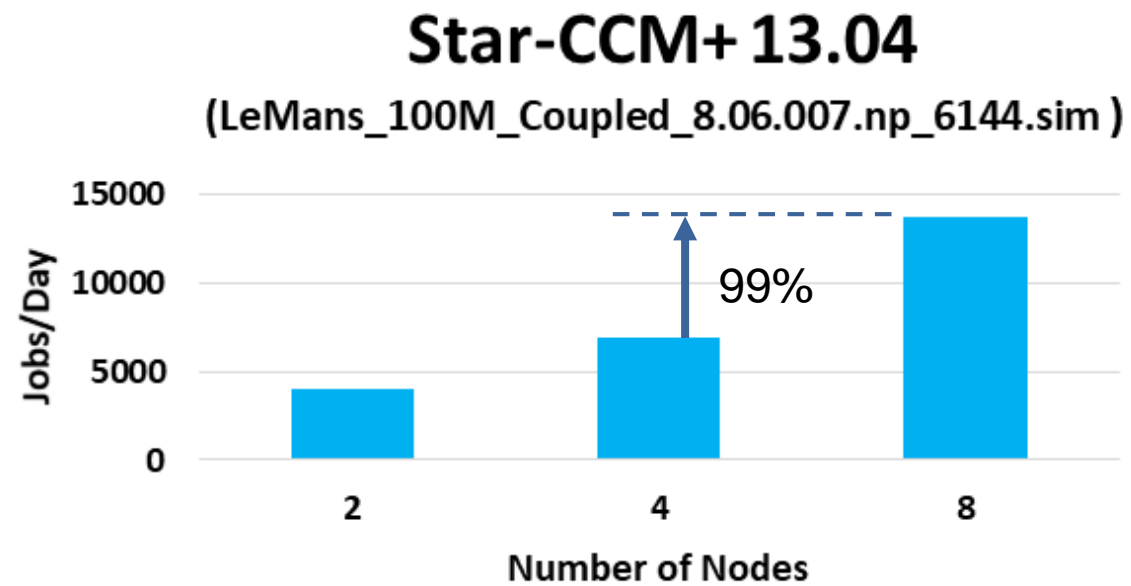
### Star-CCM+ 13.04 (EmpHydroCyclone\_30M.sim)



### Star-CCM+ 13.04 (HIMach10AoA10Sou.sim )




*Higher is better*



*Higher is better*


## Star-CCM+ 13.04 (reactor.sim)




*Higher is better*

# Star-CCM+ Application Profile on “reactor.sim”

- **34% MPI and WallClock**




- MPI\_Recv
- MPI\_Isend
- MPI\_Waitall
- MPI\_Waitany
- MPI\_Test
- MPI\_Allreduce
- MPI\_Gatherv
- MPI\_Allgatherv
- MPI\_Bcast
- MPI\_Wait
- MPI\_Alltoall
- MPI\_Irecv
- MPI\_Testall
- MPI\_Allgather
- MPI\_Comm\_split
- MPI\_Probe
- MPI\_Barrier
- MPI\_Comm\_create
- MPI\_Send
- MPI\_Comm\_dup
- MPI\_Comm\_free
- MPI\_Gather
- MPI\_Comm\_size
- MPI\_Comm\_rank
- MPI\_Comm\_group
- MPI\_Finalize
- MPI\_Init





- **Sparse Communication between the ranks, more communication between close ranks**



- Communication Pattern

Communication Event Statistics (% detail, --- error)									
	Comm Size	Buffer Size	Ncalls	Total Time	Avg Time	Min Time	Max Time	%MPI	%Wall
MPI_Recv	0	8	930184	2.667113e+03	2.867296e-03	0.000000e+00	3.640600e+00	30.94	10.59
MPI_Isend	0	8	1391731	1.440432e+03	1.034993e-03	0.000000e+00	6.953700e-02	16.71	5.72
MPI_Waitall	0	0	5628270	1.352772e+03	2.403531e-04	0.000000e+00	2.922000e+00	15.69	5.37
MPI_Waitany	0	0	34544340	9.705235e+02	2.809501e-05	0.000000e+00	5.435300e-01	11.26	3.85
MPI_Test	0	0	1102692061	5.287902e+02	4.795448e-07	0.000000e+00	2.504300e-02	6.13	2.10
MPI_Recv	0	2048	185168	1.729403e+02	9.339640e-04	0.000000e+00	2.490400e+00	2.01	0.69
MPI_Gatherv	0	28	1069	1.175768e+02	1.099876e-01	0.000000e+00	5.224300e-01	1.36	0.47
MPI_Wait	0	0	1280769	7.204506e+01	5.625141e-05	0.000000e+00	4.356800e-01	0.84	0.29
MPI_Recv	0	2560	105876	7.184717e+01	6.785973e-04	0.000000e+00	2.374800e+00	0.83	0.29
MPI_Allreduce	0	4	711168	5.420370e+01	7.621785e-05	9.059900e-06	1.632800e-02	0.63	0.22
MPI_Allreduce	0	8	980012	5.174833e+01	5.280377e-05	9.536700e-07	1.248600e-02	0.60	0.21
MPI_Isend	0	448	1064678	4.645553e+01	4.363341e-05	0.000000e+00	1.502800e-02	0.54	0.18
MPI_Isend	0	10240	152194	3.993515e+01	2.623964e-04	0.000000e+00	5.116800e-02	0.46	0.16
MPI_Isend	0	20480	33141	3.796139e+01	1.145451e-03	0.000000e+00	3.133800e-02	0.44	0.15
MPI_Bcast	0	8	96256	3.305723e+01	3.434303e-04	0.000000e+00	1.740600e-02	0.38	0.13
MPI_Allreduce	0	16	222208	3.236635e+01	1.456579e-04	8.106200e-06	1.293200e-02	0.38	0.13
MPI_Alltoall	0	4	39936	3.182579e+01	7.969197e-04	7.605600e-05	6.228100e-02	0.37	0.13
MPI_Bcast	0	1	116736	3.083769e+01	2.641661e-04	0.000000e+00	1.031100e-02	0.36	0.12
MPI_Recv	0	4096	111592	2.715308e+01	2.433246e-04	0.000000e+00	2.520200e+00	0.31	0.11
MPI_Recv	0	512	64179	2.640617e+01	4.114457e-04	0.000000e+00	2.766400e+00	0.31	0.10
MPI_Isend	0	384	1204437	2.515343e+01	2.088397e-05	0.000000e+00	1.641600e-02	0.29	0.10
MPI_Recv	0	5120	93042	2.301502e+01	2.473616e-04	0.000000e+00	3.622000e+00	0.27	0.09
MPI_Isend	0	16384	71641	2.177818e+01	3.039905e-04	0.000000e+00	7.943400e-02	0.25	0.09
MPI_Recv	0	6144	77613	2.176994e+01	2.804934e-04	0.000000e+00	3.583600e+00	0.25	0.09
MPI_Isend	0	256	1761913	2.050793e+01	1.163958e-05	0.000000e+00	1.596000e-02	0.24	0.08
MPI_Isend	0	512	2455865	1.896468e+01	7.722200e-06	0.000000e+00	1.489000e-02	0.22	0.08
MPI_Recv	0	8192	91164	1.766754e+01	1.937995e-04	0.000000e+00	2.164900e+00	0.20	0.07
MPI_Recv	0	3072	60215	1.552832e+01	2.578813e-04	0.000000e+00	2.445600e+00	0.18	0.06

- **Star-CCM+ performance testing over AMD EYPC based platform**
  - An average of 98% scaling was achieved from 4 to 8 nodes among all 5 different input files
- **Star-CCM+ MPI profiling on “reactor”**
  - MPI communication accounts for 34% of overall wall clock time at 8 nodes
  - MPI\_Recv is 33% of MPI, MPI\_Isend is 30% of MPI and MPI\_Waitall is 20% of MPI
  - Most communication is sparse, among all nodes. Some more communication between near ranks



# Thank You

