SPARC INTEL COMPILER

GCC (64 cores. 1 node MUMPS)

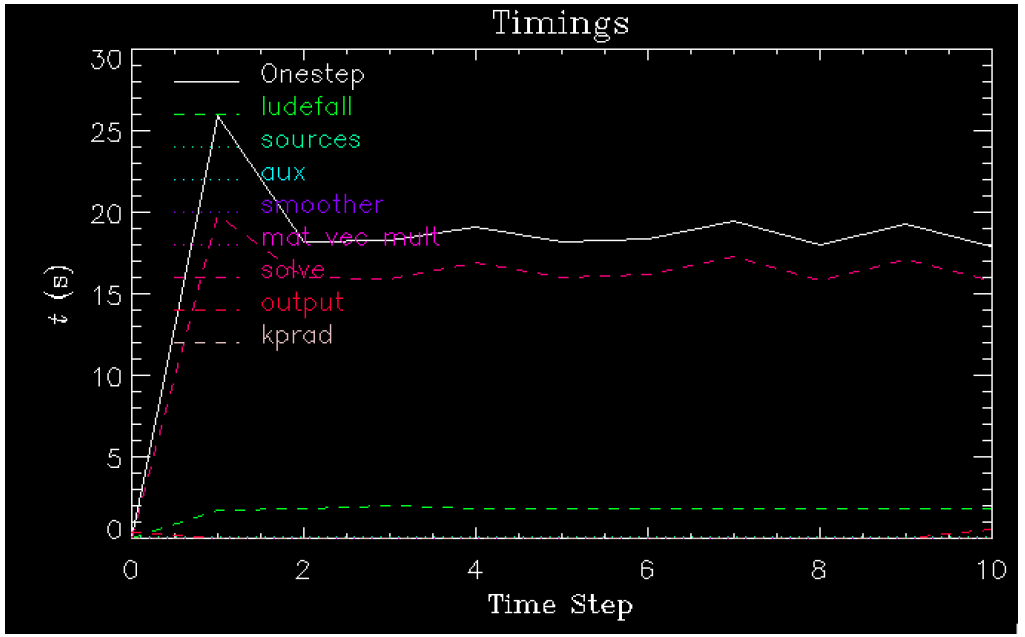

INTEL (64 cores. 1 node MUMPS)

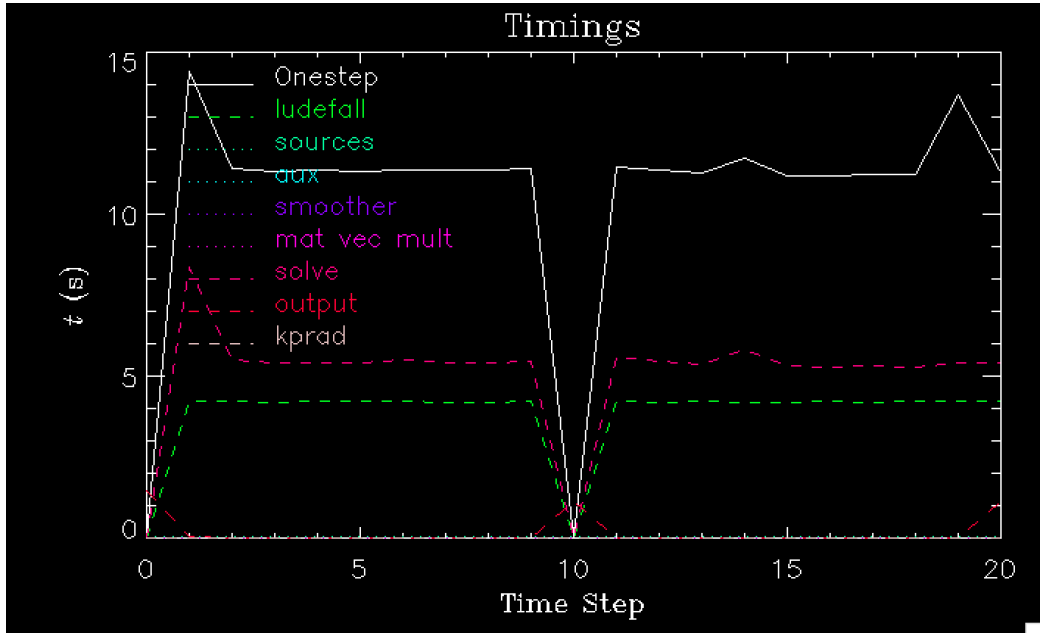

- **”solve” metric (red) substantially improved with intel compilers**
- **“ludefall” metric (green) worsened**
- **Overall performance “Onestep” (white) improved**
- **Both cases launched from local directory.**
  INTEL case from ”scratch” does not show improvements

SPARC INTEL compilation

Comparison with other systems

Sparc INTEL (64 cores. 1 node MUMPS)

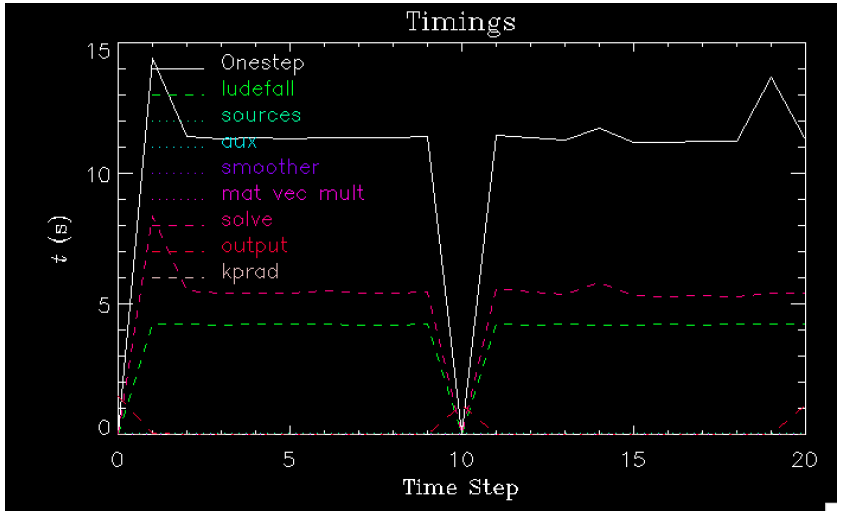

Stellar (64 cores. 1 node MUMPS)

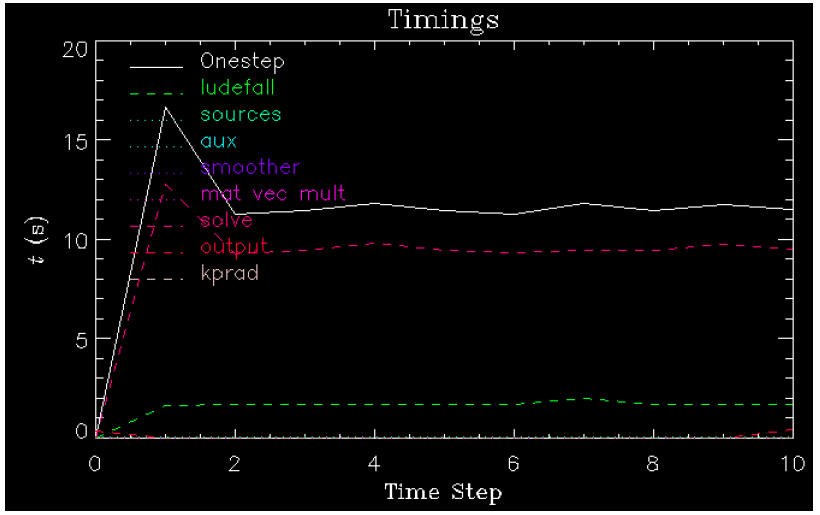

Perlmutter (64 cores. 1 node. MUMPs

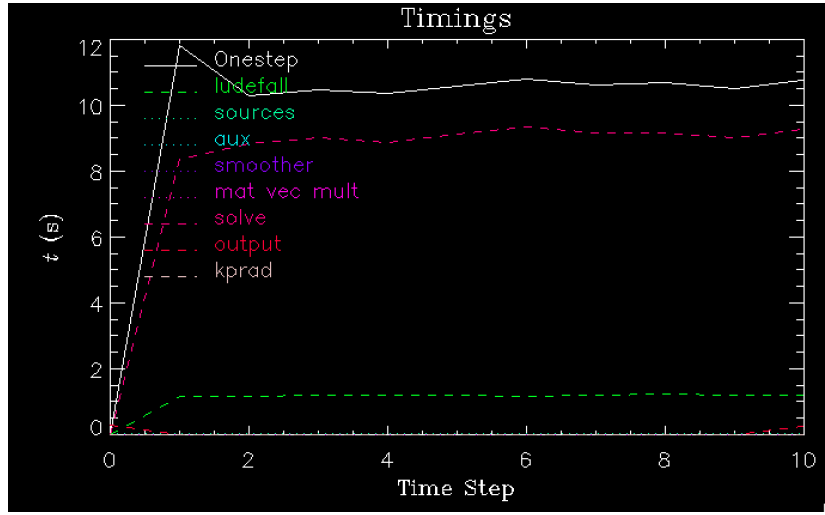

- **”solve” metric better than other systems**
- **“ludefall” metric worse**
- **Overall performance “Onestep” similar across systems**

Sparc Intel compilation: scalability

32 CORES (1 NODE)

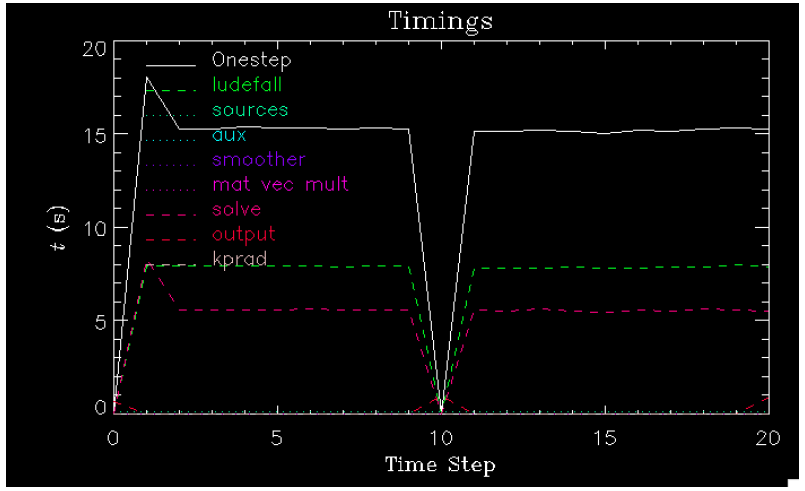

64 CORES (1 NODE)

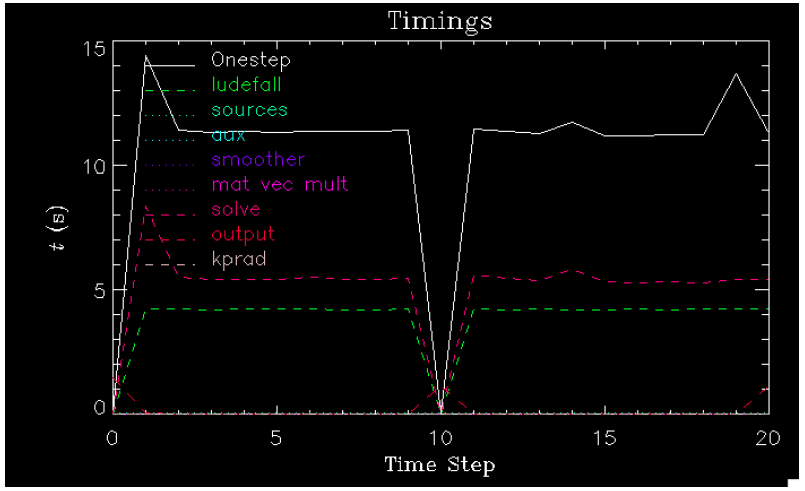

128 CORES (2 NODES)

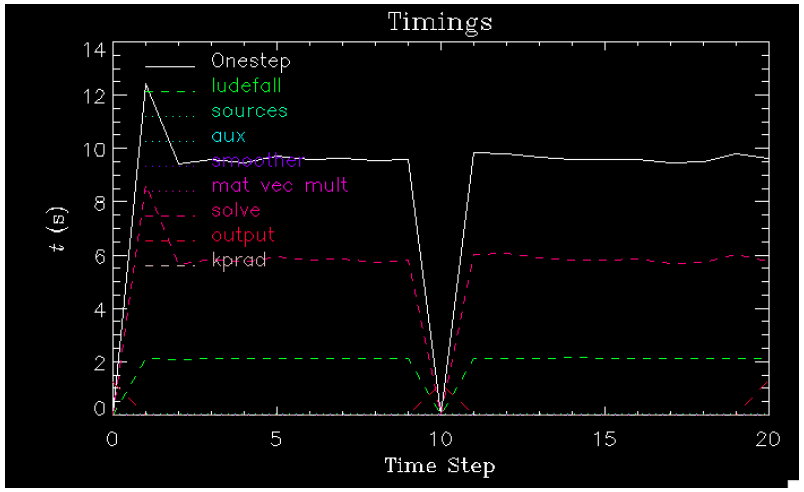

192 CORES (3 NODES)

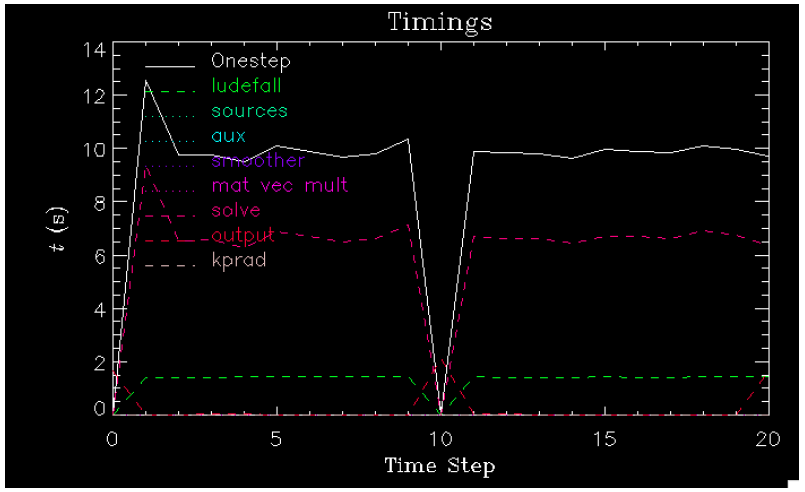

- **‘solve’ metric does NOT scale**
- **‘ludefall’ metric scales well**
- **Due to the significant time consumed by ‘ludefall’, the overall performance ‘onestep’ improves.**

SPARC INTEL Compilation: 3D Simulation scaling

64 mesh partitions x 4 planes = 256 cores
(SPARC: 4 nodes)

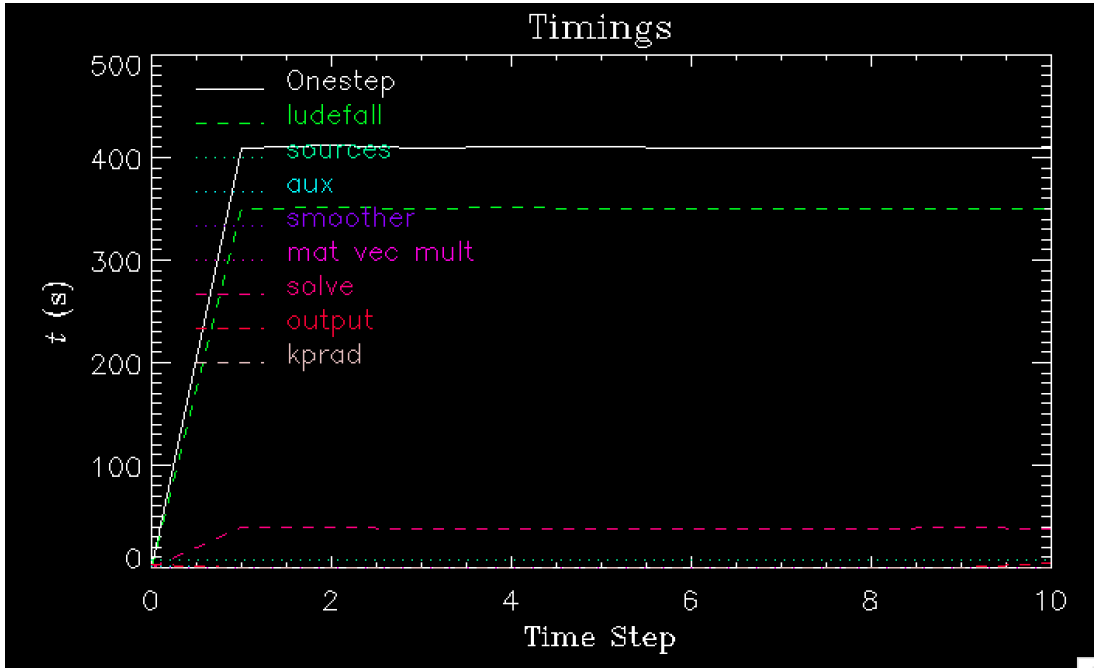

128 mesh partitions x 4 planes = 512 cores
(SPARC: 8 nodes)

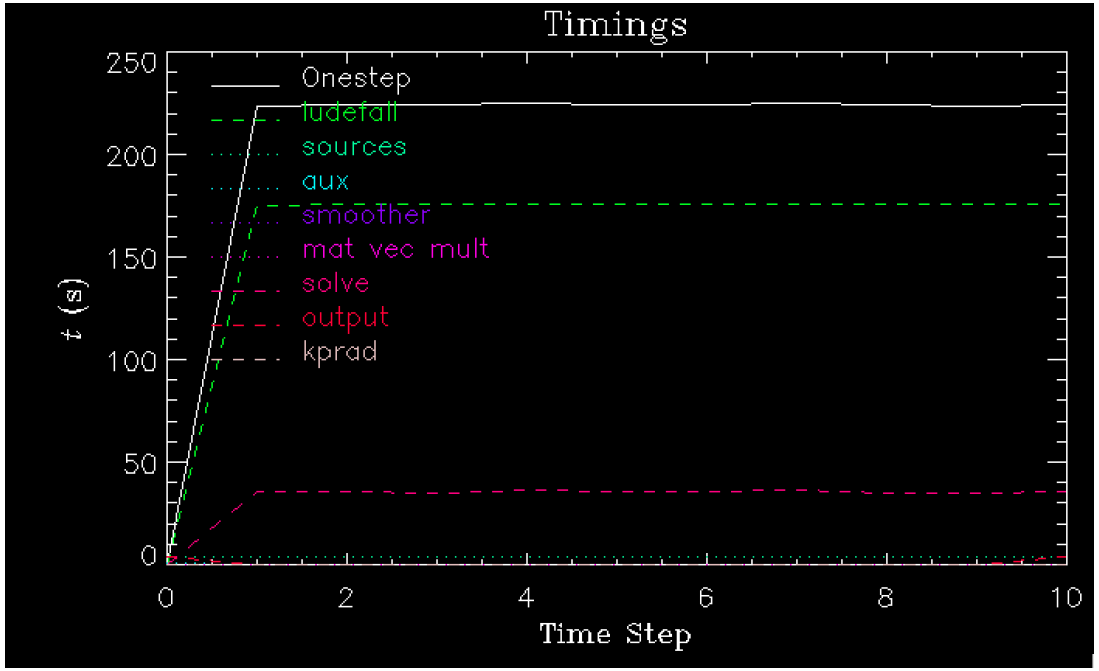

- **'solve' DOES NOT scale (~40 sec / time step in both cases)**
- **'ludefall' DOES SCALE from 350 sec down to 175 sec per time step (x2)**
- **'onestep' metric goes from 410 sec down to 220 sec per time step (x1.9)**
  onestep is dominated by ludefall.

Princeton STELLAR: 3D Simulation Scaling

64 mesh partitions x 4 planes = 256 cores
(STELLAR: 3 nodes)

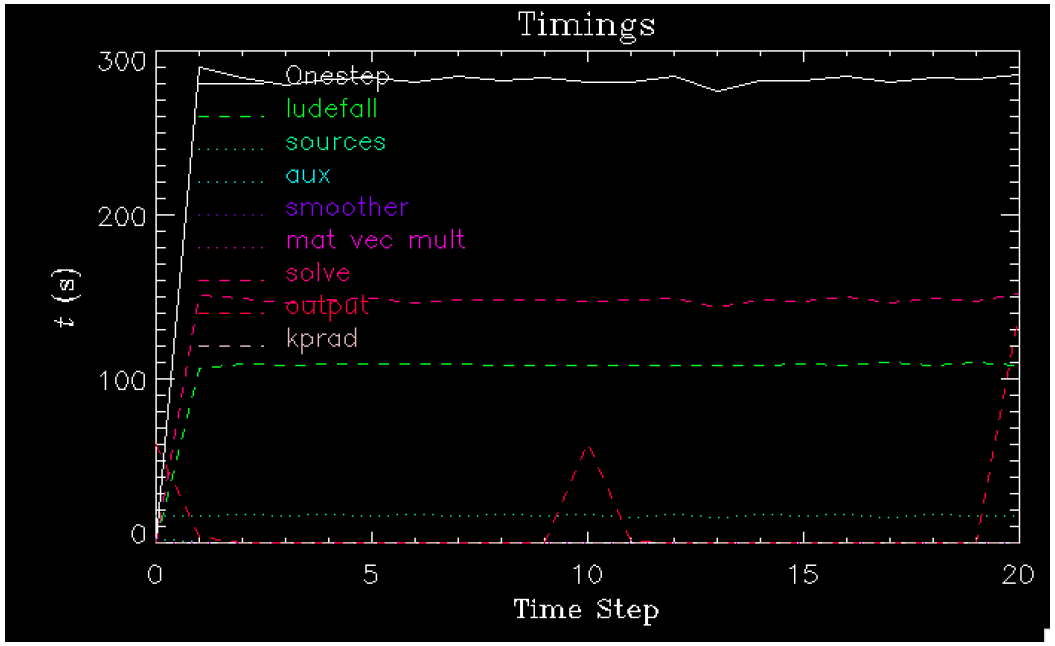

STELLAR: 96 mesh partitions x 4 planes = 384 cores (4 nodes)

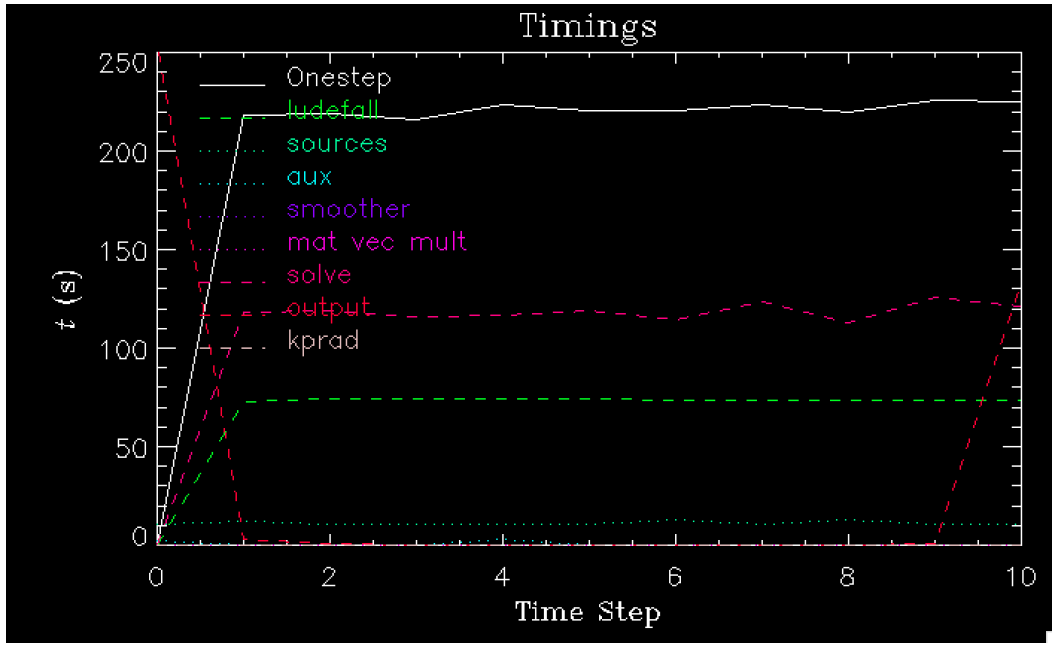

- **‘solve’ goes from 150 sec down to 120 sec per time step (x1.25)**
- **and ‘ludefall’ goes from 110 sec down to 70 sec per time step (x1.5)**
- **‘onestep’ metric goes from 290 sec down to 220 sec per time step (x1.3)**

In good agreement with the increase in the number of cores