ANSYS HPC Seminar Series

Mechanical

• Prepared and presented by Alex Pickard

apickard@simutechgroup.ca

Who Are We

- 120+ Employees
- 12 Local Offices
- 2000+ Customers
- 80% Engineering Staff with Advanced degree

Today’s Adenda

- Introduction to Mechanical HPC, solvers, and licenses
- Computer platform recommendations (CPU models & features, RAM quantity, memory bandwidth, storage, networking)
- Hardware recommendations for users (solver and preprocessor)
- Performance diagnostic information from reading solver logs
- Simulation best practices for performance and scalability
- Sizing resource request for your model
- Intro to clusters and Job Schedulers
- Workflow recommendations based on analysis type
- Demos / Workshops
  - Job submission techniques
  - Multi-step simulation (3D Printing)
  - Improving performance and scalability of contacts

What is HPC?

High-performance computing is the use of parallel processing techniques for solving complex computational problems.

It especially refers to using multiple computers to work together on a single problem (clusters). It does not necessarily mean working on a single solution.

What is HPC?

Having HPC capability increases throughput:

- Faster results
- More design iteration
- Hit hard deadlines
- Greater engineering efficiency
- Parametric analysis & optimization

HPC enables more thorough design and analysis on a tighter deadline.

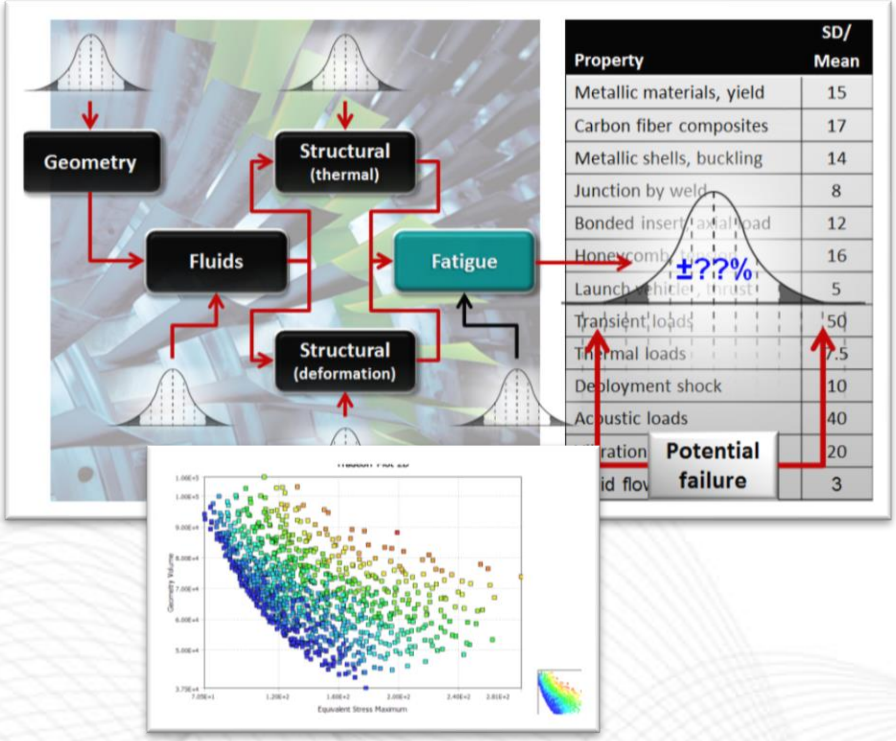

What is HPC?

Having HPC capability increases capability:

- More complex models
- More physics
- Less risky simplifications
- Greater detail
- System level analysis
- Discovery of new insight

HPC unlocks new capability within the ANSYS products your already have.

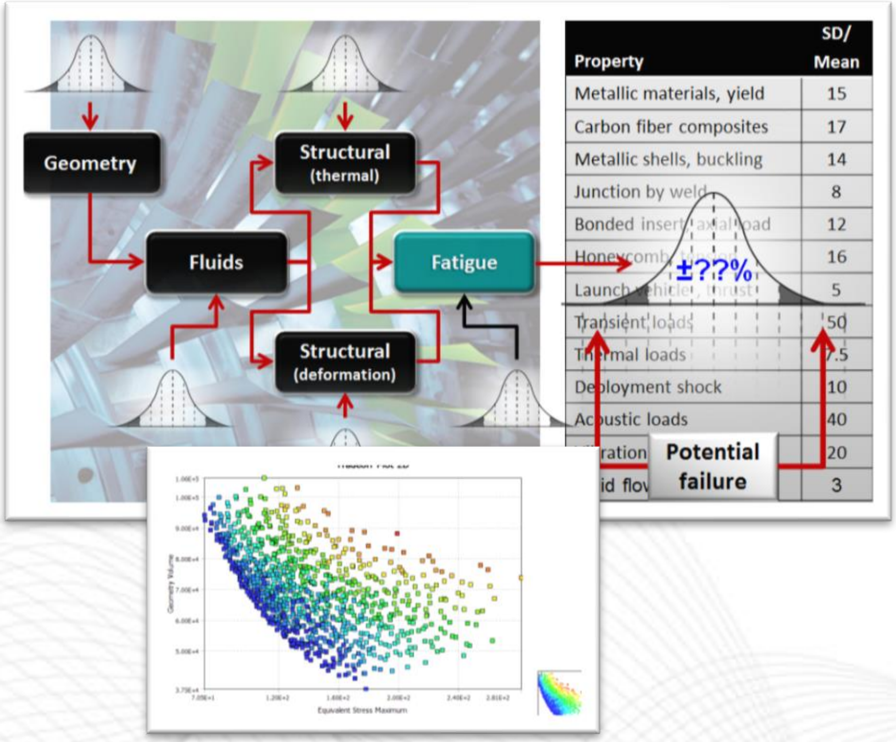

ANSYS Mechanical

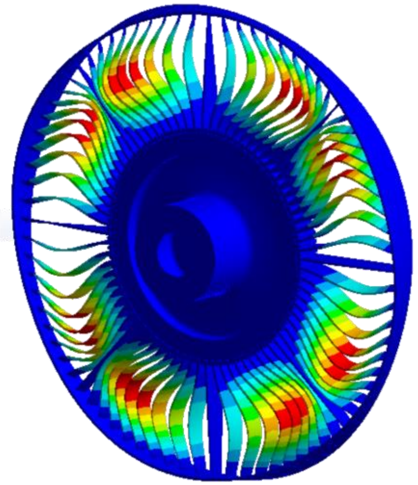

- The ANSYS Mechanical suite is a premium set of finite element analysis tools for a wide variety of problems:
  - Stress Analysis
  - Fatigue
  - Composites
  - Vibration and dynamics
  - Explicit (impact)
  - Acoustics

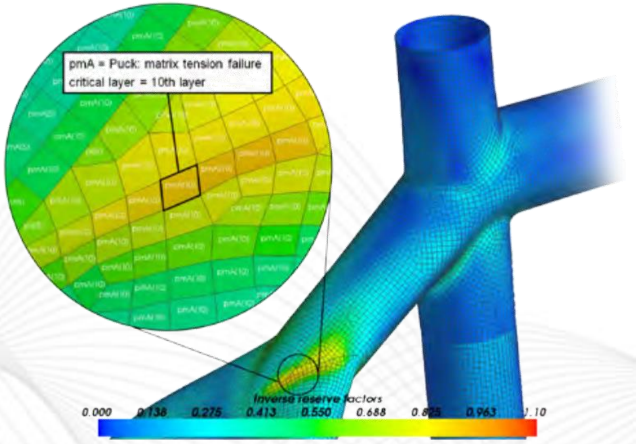

ANSYS Mechanical Solvers

- The ANSYS Mechanical solvers are a set of mathematical approaches to solving the problem as input:
  - Sparse Matrix Direct Solver (direct elimination)
  - PCG Solver (iterative)
  - JCG Solver (iterative)
  - ICCG Solver (iterative, Shared Memory Only)
  - QMR Solver (iterative , Shared Memory Only)

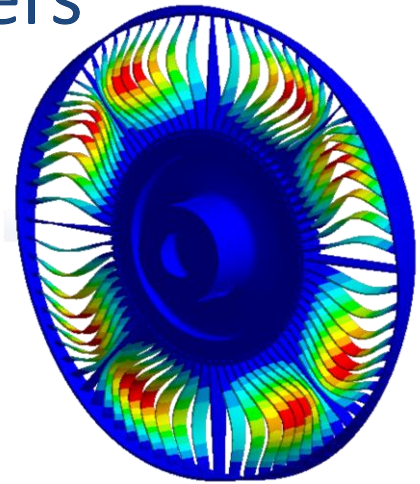

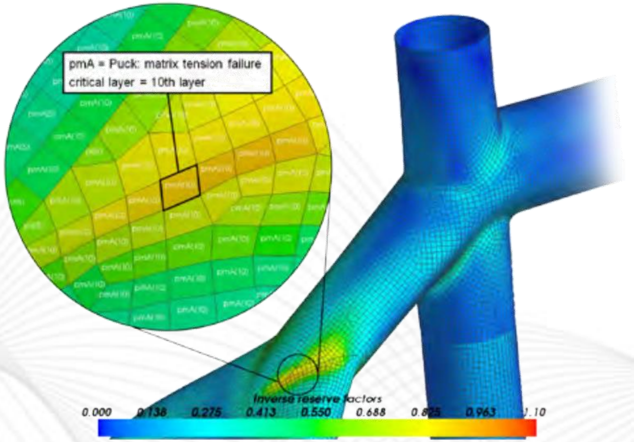

Distributed Memory Mode

- Distributed solvers are standard across many ANSYS products.
- They scale and perform better than shared memory solvers.
- They enable each CPU core to have it’s own solver process and work on and independent chunk of the problem
- Requires substantial coordination and communication between processes.
- Enables problems to be run across multiple computers.
- Many technologies are difficult to distribute across independent tasks:
  - Eigensolver, nonlinear adaptive remeshing

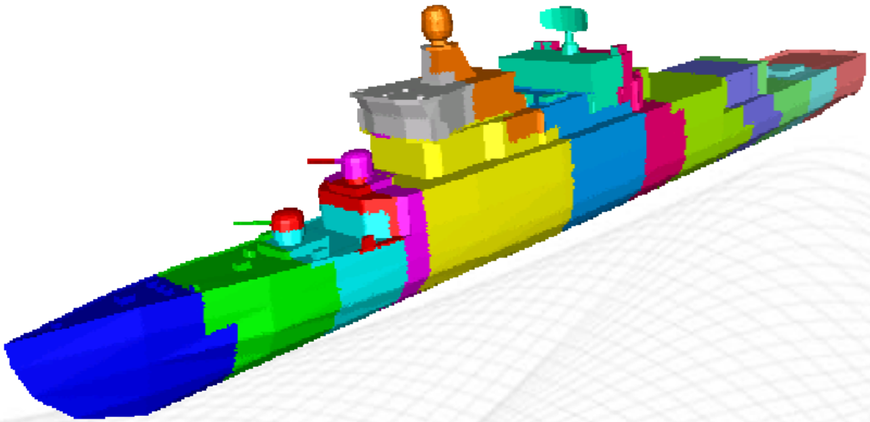

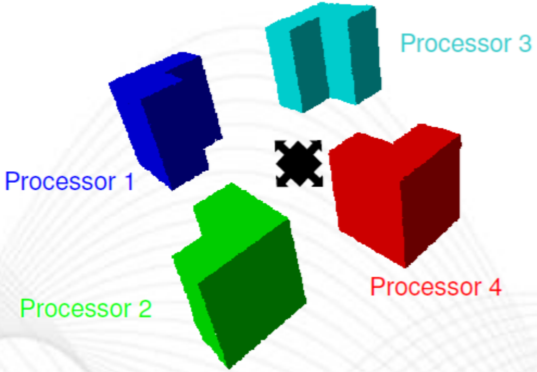

Distributed Memory Mode

- Example of shared memory (SMP) vs. distributed memory (DMP) scaling:

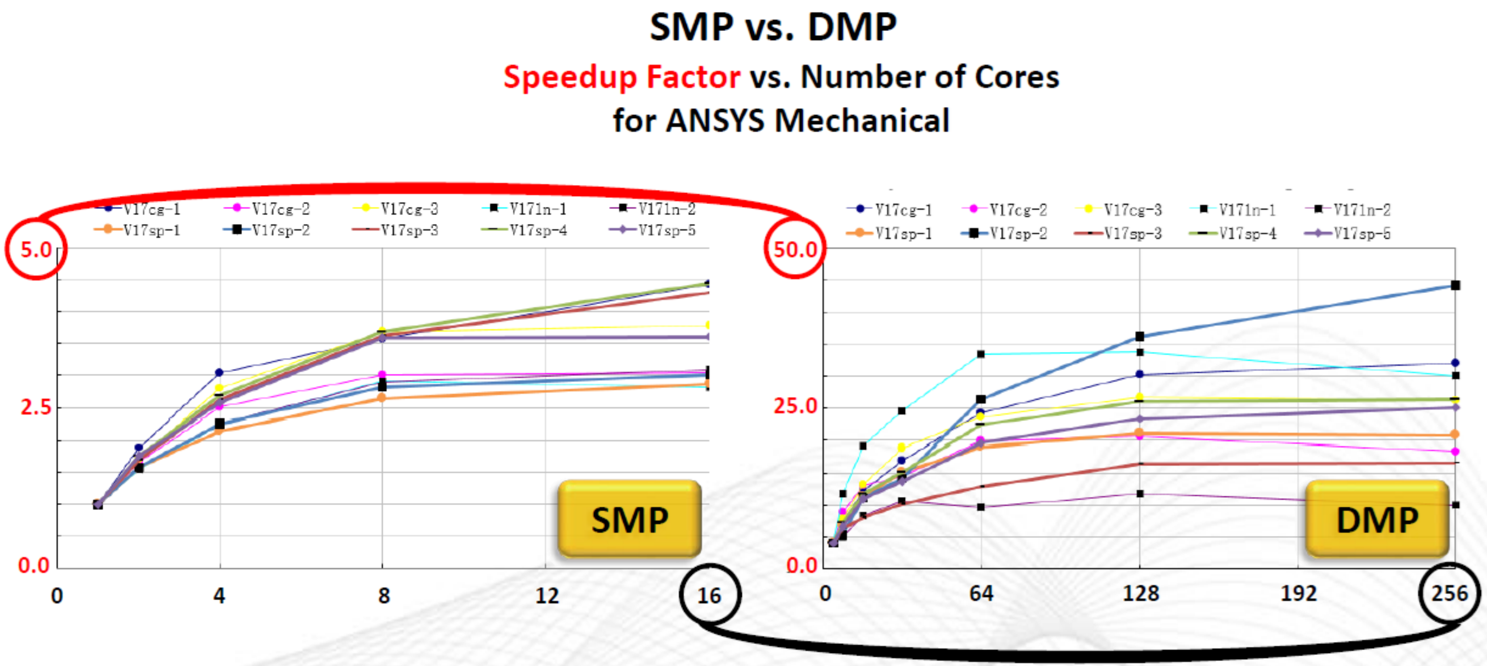

ANSYS Mechanical Solvers

- **Sparse Matrix Direct Solver (direct elimination)**
  - When robustness and solution speed are required (nonlinear analysis)
  - For linear analysis where iterative solvers are slow to converge (ill-conditioned matrices)
- **PCG Solver (iterative)**
  - Reduces disk I/O & RAM requirement relative to sparse solver.
  - Best for large models with solid elements and fine meshes.
  - Most robust iterative solver in ANSYS. Much faster than JCG, but uses more RAM.
- **JCG Solver (iterative)**
  - Best for single field problems - (thermal, magnetics, acoustics, and multiphysics).
  - Uses a fast but simple preconditioner with minimal memory requirement.
  - Not as robust as PCG solver.
- ICCG Solver (iterative, Shared Memory Only)
  - More sophisticated preconditioner than JCG.
  - Best for more difficult problems where JCG fails, such as unsymmetric thermal analyses.
- QMR Solver (iterative , Shared Memory Only)
  - Used for full harmonic analyses.

Solver Memory Requirements

Solver	Ideal Model Size	Memory Usage SMP	Memory Usage DMP	Disk usage
Direct	100k -10M DOF	In RAM: 10+GB/MDOF Paging to Disk: 1GB/MDOF	In RAM: 15+GB/MDOF Paging to Disk: 1.5GB/MDOF	In RAM: 1.5GB/MDOF Paging to Disk: 15+GB/MDOF
PCG	500k-20M+ DOF	0.3 GB/MDOF with MSAVE,ON 1 GB/MDOF without MSAVE	2 GB/MDOF	0.5 GB/MDOF
JCG	500k-20M+ DOF	0.5 GB/MDOF	0.5 GB/MDOF	0.5 GB/MDOF

- DOF: Degree of Freedom in the matrix
  - Solid structural elements: 3 DOF / node
  - Beam and Shell structural elements: 6 DOF / node
  - Thermal Elements: 1 DOF / node
- Using more cores increases RAM usage. If you don’t have enough try dialing back the core request.

ANSYS Academic & HPC Licence Changes

- ANSYS HPC licencing has changed multiple times in the last 5 years.
- Overall trend is to become (slightly) simpler.
- As of 2021 Academic and Commercial licencing is identical:
  - Previously there were dedicated research licenses that provided 16 cores standard.
  - Now academic and campus licence bundles contain **commercial** licenses.
  - Campus bundles contain more HPC to compensate, more than makes up for the change.
    - Old 10 Research bundle had 10 x 16 core solvers + 64 HPC (no solve bigger than 80 cores)
    - New 10 Research bundle has 10 x 4 core solvers + 180 HPC (allows up to 184 cores in one solve)
- Sovlers enable 4 cores standard + HPC licence to add individual cores.
- HPC Packs are also available which have an exponential effect:
  - First pack **triples** the allowed cores: 4 -> 12 = 3X
  - Second pack **triples** it again: 12 -> 36 = 3X
  - Next ones is nearly **quadruple**: 36 -> 132 = ~3.7X
  - Note: Uncommon in academia… 132 -> 516 = ~3.9X

Packs	Added Cores	Total Cores
1	8	12
2	32	36
3	128	128
4	512	516
5	2048	2052

CPU Instructions: AVX

- The direct solver uses AVX (Advanced Vector Extensions) to solve the matrix.
- AVX performs a single instruction on multiple pieces of data (SIMD).
- Increases parallelism within the CPU core, and thus performance.
- The latest AVX incarnation (AVX-512) is enabled in ANSYS 18.2.
- The AVX functionality is provided by the Intel Math Kernel Library (Intel MKL).
- We want AVX2+

- R18 Benchmark set (DMP)
- Used geometric mean values for each class of benchmarks
- Used 1, 2, 4, 8, 16, & 32 cores
- 2 Intel Xeon Gold 6148 (2.4 GHz, 40 cores total), 192 GB RAM, Linux CentOS 7.3

	Iterative Solver Benchmarks	Direct Solver Benchmarks
R18.1	557 sec	474 sec
R18.2	537 sec	319 sec

R18.2 performs over 30% faster than R18.1 on Skylake systems

ANSYS HPC Technology Advancement

- AVX or Advanced Vector Extensions:
  - Pack more operations in a single clock cycle
  - Has been progressively improved over CPU generations

	Microarchitecture	Instruction Set	SP FLOPs / cycle	DP FLOPs / cycle
2017	Skylake	Intel® AVX-512 & FMA	64	32
2014	Haswell / Broadwell	Intel AVX2 & FMA	32	16
2012	Sandybridge	Intel AVX (256b)	16	8
2010	Nehalem	SSE (128b)	8	4

Regarding MKL and AMD

- The AVX functions used to directly solve the matrix are enabled via the Intel MKL.
- Intel MKL checks for “Genuine Intel” description on the CPU before enabling AVX.
- ANSYS versions prior to 2020R1 ran quite poorly on AMD CPUs if a workaround is not used.
- Intel committed to improving support for AMD in their Math Library, so 2020R1+ performance is greatly improved, but the workaround is still better.
- Workaround is expected to be disabled in the MKL version shipping 2021R2, but AMD support is also expected to continue improving.
- Workaround: Environment Variable forces AVX support
  MKL_DEBUG_CPU_TYPE=5
- AMD AVX2 support:
  - Half speed on Ryzen 1000 & 2000 & EPYC Naples (1 AVX instruction per 2 clock cycles)
  - Full speed on Ryzen 3000-5000 & EPYC Rome

AVX Support by Model

- AMD CPUS:
  - No AVX-512 support.
  - Substandard treatment by Intel MKL
  - Otherwise really good!
- Intel CPUs:
  - Some have AVX-512: HEDT LGA2066, Xeon W, Xeon Scalable Bronze-Platinum
    - Xeon Gold 6000+ has 2 AVX-512 Exection units per core
    - Xeon Bronze, Silver, and Gold 5000 have only 1 AVX-512 unit enabled.
  - Many only have AVX2: Consumer Laptop and Desktop CPUs, Xeon E series
- It’s not all down to AVX.
  - AMD CPUs seem superior for iterative solvers
  - AMD EPYC CPUs have substantial memory bandwidth

Memory Bandwidth

- Simulation software tends to do simple math on large pools of data.
- Data needs to fetched quickly, processed, and then the results stored quickly.
- The ability to fetch and store data is limited by the memory bandwidth and data cache on a system, which is inherent to the platform and CPU.
- Comparison of current server CPU models available:

Model	**Intel Xeon Scalable (per socket)**	**AMD EPYC (per socket)**
Memory Channels	6 x DDR4-2933 MHz	8 x DDR4-3200 MHz
L3 Data Cache	Up to 38.5 MB	Up to 256 MB
Cores	Up to 28	Up to 64

- As core count gets higher the bandwidth advantages exceed the core efficiency

Memory Bandwidth

- Be sure to populate all memory channels (not necessarily all slots)
- Aim between between 2-4 Cores / memory channel
- https://simutechgroup.com/maximizing-memory-performance-for-ansys-simulations/

8 Core Test (CFX)

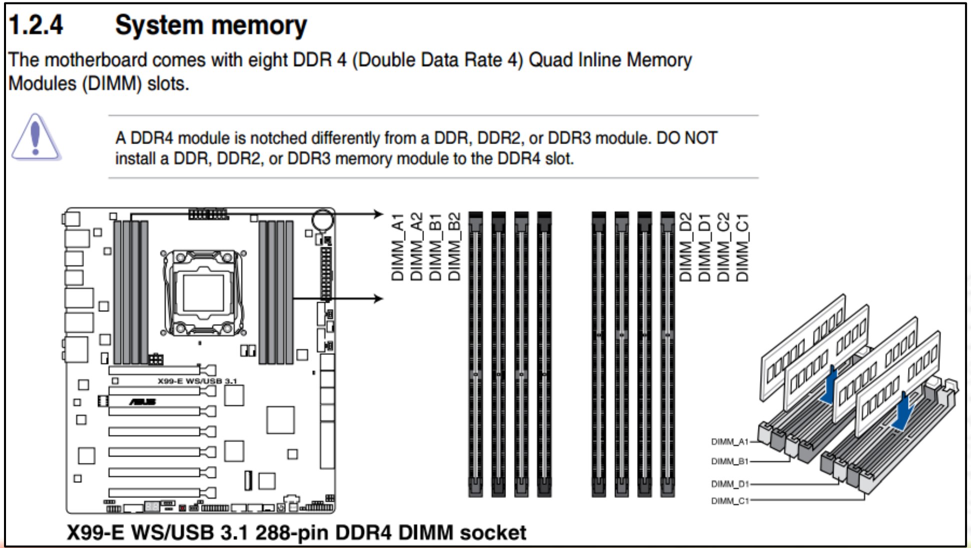
1.2.4 System memory

The motherboard comes with eight DDR 4 (Double Data Rate 4) Quad Inline Memory Modules (DIMM) slots.

A DDR4 module is notched differently from a DDR, DDR2, or DDR3 module. DO NOT install a DDR, DDR2, or DDR3 memory module to the DDR4 slot.

X99-E WS/USB 3.1 288-pin DDR4 DIMM socket

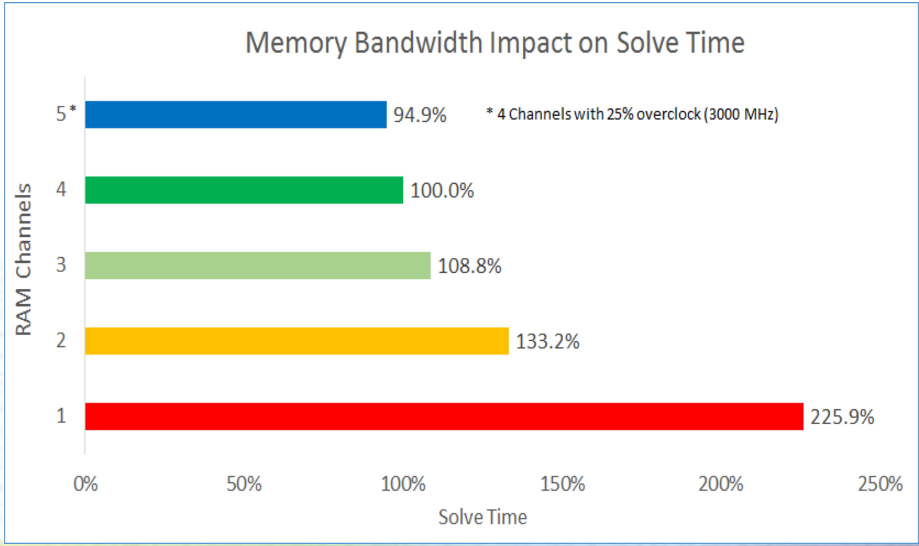

Which CPU Then?

- The direct solver is really important for Mechanical; iterative solvers can’t solve all problems.
- Generally recommend Intel due to better MKL support and core efficiency at low to modest core counts.
- If licencing doesn’t matter and looking at high core count then consider AMD.
- Either for laptops
- Focus on frequency for preprocessing

Cores / Node	Iterative	Direct
4-8	Ryzen 5800X i9-10900k	i9-10900X Xeon W
12-16	Threadripper 3960X Ryzen 5950X	i9-10920X / i9-10980XE Xeon W Threadripper 3960X
24+	Dual AMD EPYC	Dual Xeon Gold (like 6242R) Dual AMD EPYC

Dual EPYC 7502 vs Dual Xeon 6242R
64 cores vs 40 cores

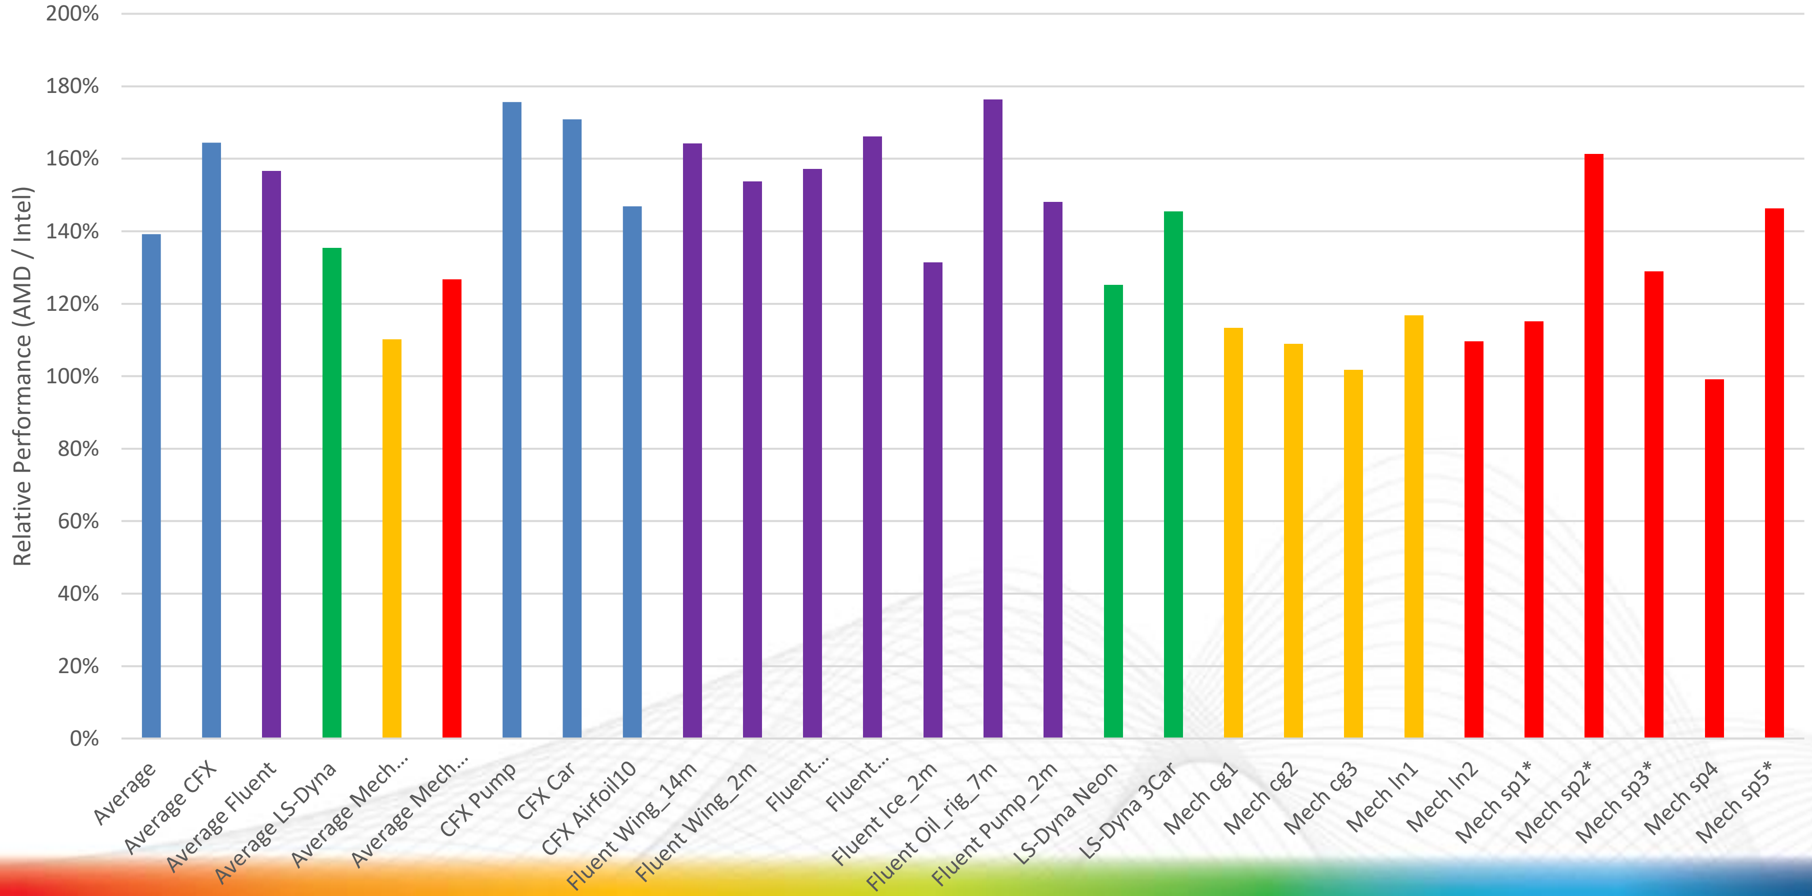

Dual EPYC 7502 vs Dual Xeon 6242R
36 cores vs 36 cores

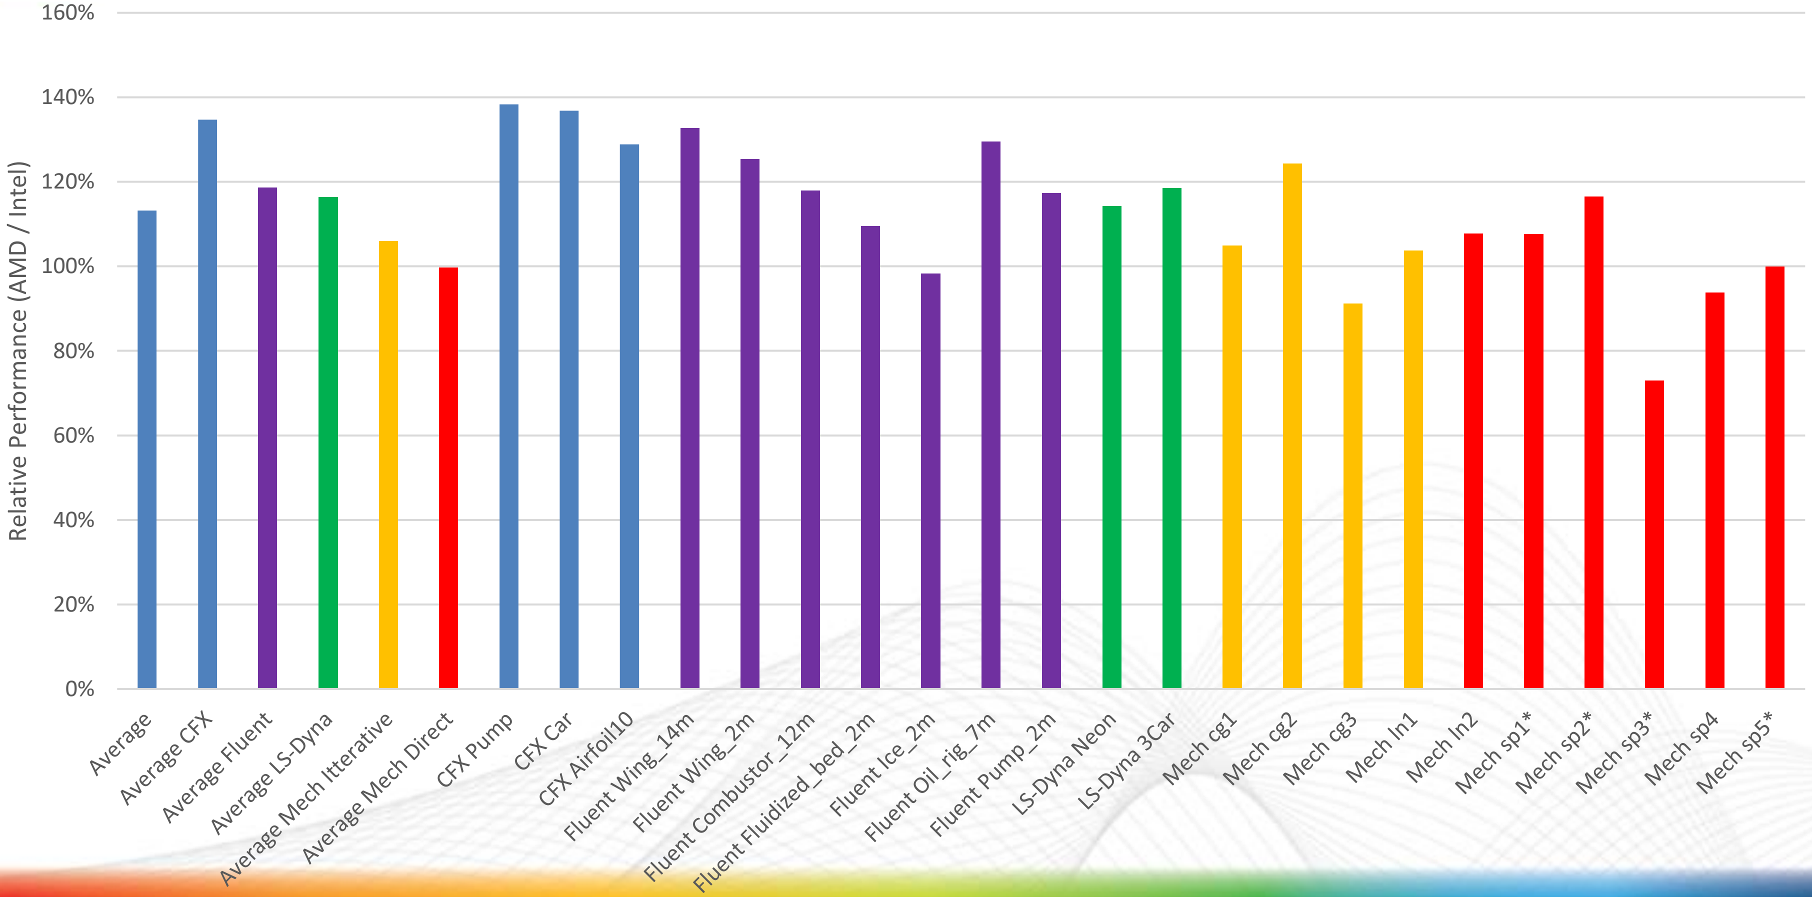

Dual EPYC 7502 vs Dual Xeon 6242R
36 cores vs 36 cores

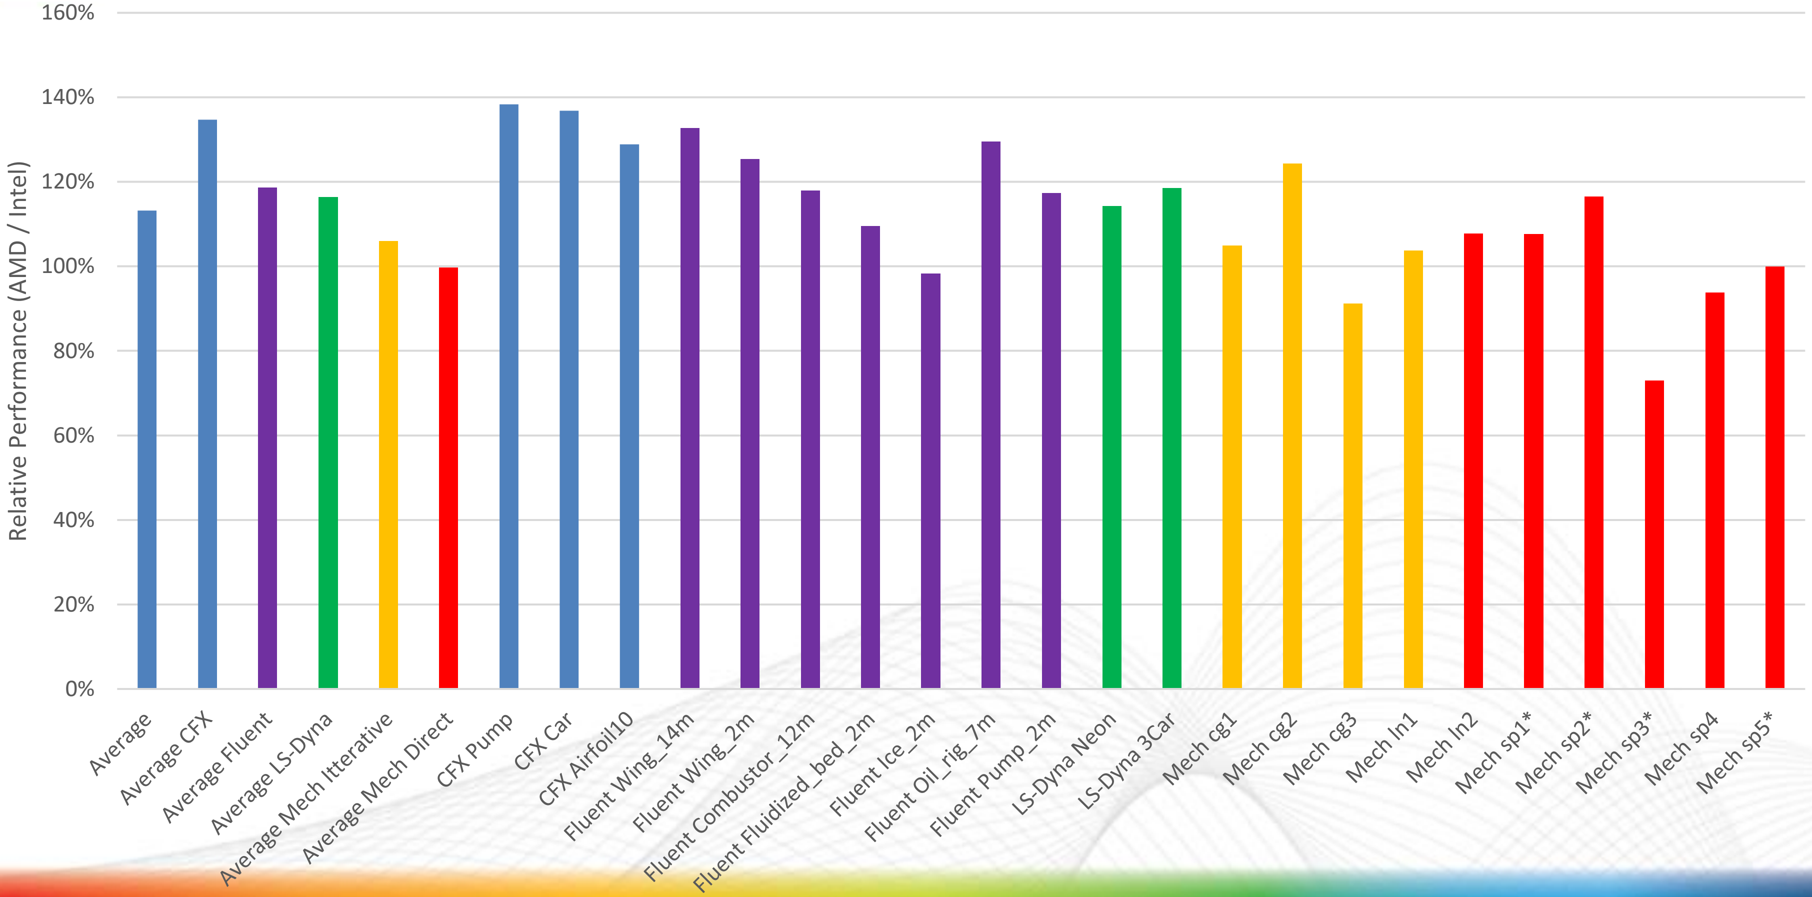

Threadripper 3960X vs i9-9960X
12 cores vs 12 cores

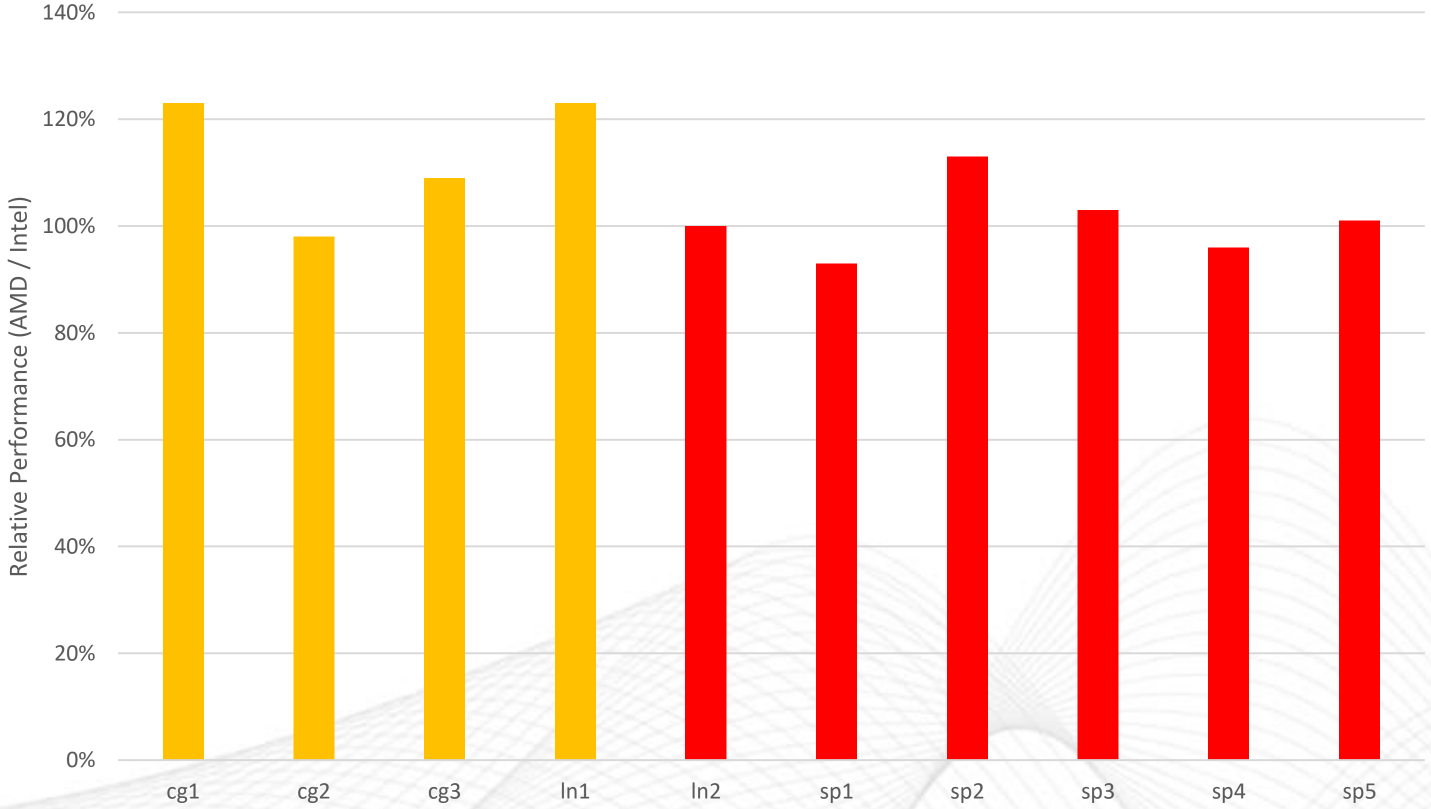

Threadripper 3960X vs i9-9960X
24 cores vs 16 cores

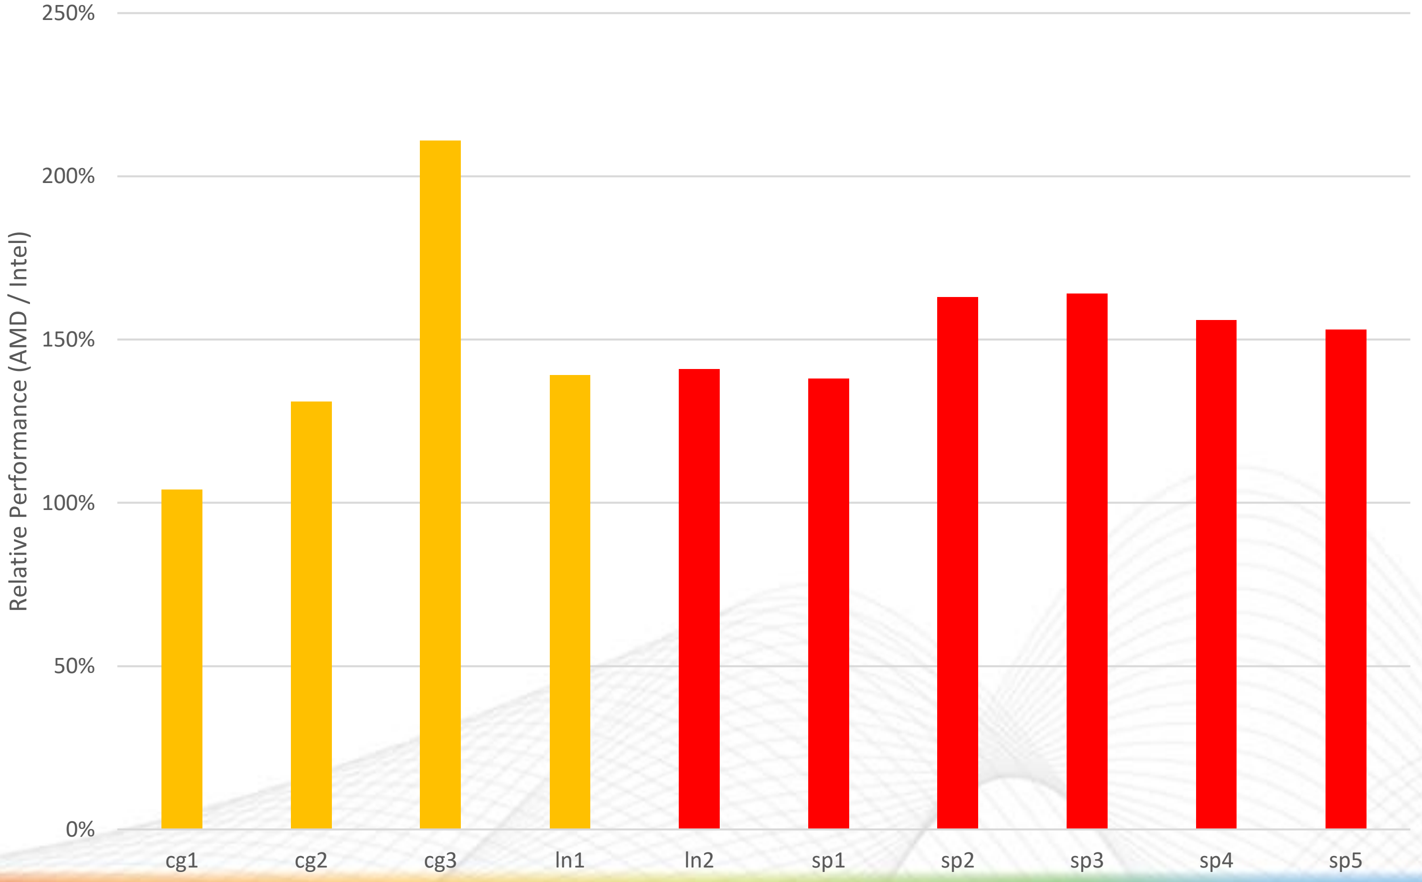

Regarding RAM Sizing

- Giving a memory recommendation is difficult because it is very usage dependent, driven by fidelity requirements.
- Extra RAM beyond what is needed to solve in RAM is not helpful.
- Sizing for 100% of use cases is impractical, aim to solve in RAM 90%+ of cases
- Largely comes down to economics and appetite, RAM is fairly cheap.
- I personally regularly work in the 2-5 Million node range which means I need roughly 128-256 GB RAM for most of my work, but often require 512 GB for single solve and on rare occasion more.
- We use a cluster to scale our RAM, more than scale our cores.
  - Need more RAM: request more nodes.
- Use NVMe SSDs to maximize performance when not fitting in RAM.

RAM Management

- The performance impact of solving out of RAM can be substantial, sometimes it is possible to make adjustments to avoid it.
  - Switching to iterative
  - Remeshing
  - Reduce bandwidth of matrix (scoping of boundary conditions)
  - Contact scoping / trimming
  - Force in-core if very close or mis-predicting DSPOPTION,,INCORE
  - Use less cores (RAM usage increases with cores for distributed solves, sometimes by a lot)
  - Switch to shared memory (equivalent to 1 core, but faster)
  - Submodeling
- Sometimes out-of-core (RAM) isn't that painful. Depends on storage

GPUs

- Dedicated GPU is strongly recommended for GUI usage. (Quadro / Firepro)
- GPU’s for compute are not really recommended
  - Very expensive
  - Inconsistent benefit
  - Less benefit as more cores are available
  - Main use case: reduce licencing costs (they only count as 1 core)

Networking

- Networking nodes together in Mechanical works very well.
- Each node adds more cache, memory bandwidth, memory quantity, cores, etc.
- High speed interconnect is required (40 gbps+).
- RDMA communication is effectively required to see proper gains.
    - Traditionally infiniband was recommended, but ethernet has this too.
    - ~2 microsecond vs 30 microsecond latency, higher bandwidth, less overhead
- Linux easier to implement, but Windows 10 and Server also work.

Job Schedulers & RSM

- Remote Solve Manager (RSM) is a background solving feature that allows solves to happen as a background task on your computer or others on your network.
- RSM can also integrate with the following Job Schedulers:
  - ANSYS RSM Cluster (ARC)
    - Windows and Linux, Free from ANSYS
  - PBS Pro (Linux)
  - Torque (Linux, very similar to PBS, just had official support dropped but still works)
  - Platform LSF (Linux)
  - SLURM (Linux, newly Supported)
  - Univa Grid Engine (Linux)
  - Windows HPC Server (Windows Server)
- Mechanical and Workbench can both be batch solved manually, allowing any scheduler.

Job Submission Techniques - RSM

- RSM is by far the preferred method to use remote solve resources with Mechanical.
- Using RSM you can:
  - Submit from within Mechanical
  - Submit Mechanical System from Workbench
  - Submit whole project from Workbench

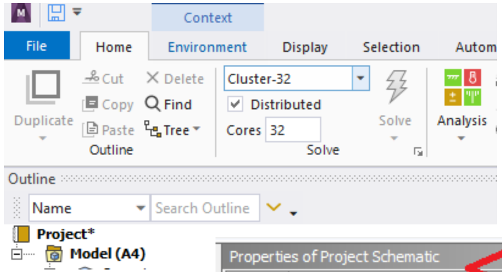

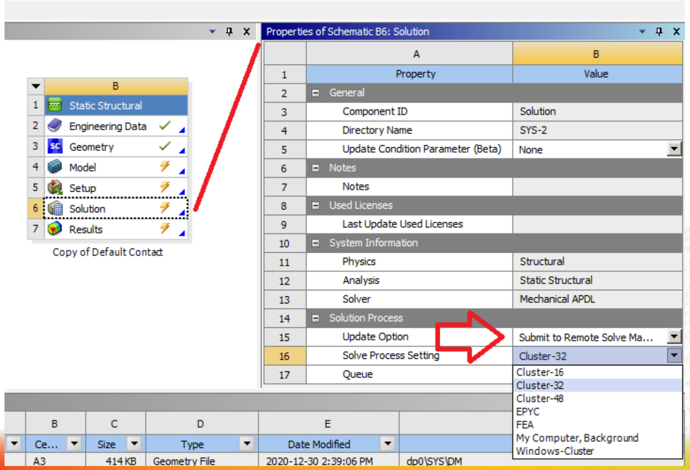

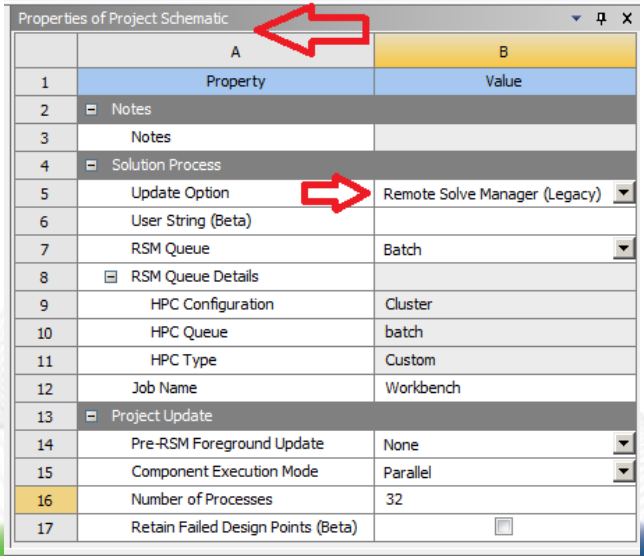

Workbench Job Submission via RSM

- Submitting whole workbench project has significant benefit for interdependent systems.
- Normally System A work be solved, the results retrieved, System B would map those results, then be submitted, solved, and retrieved.
  - Lots of file transfer and user input required.
- Workbench project update allows a single submission that updates the whole project.

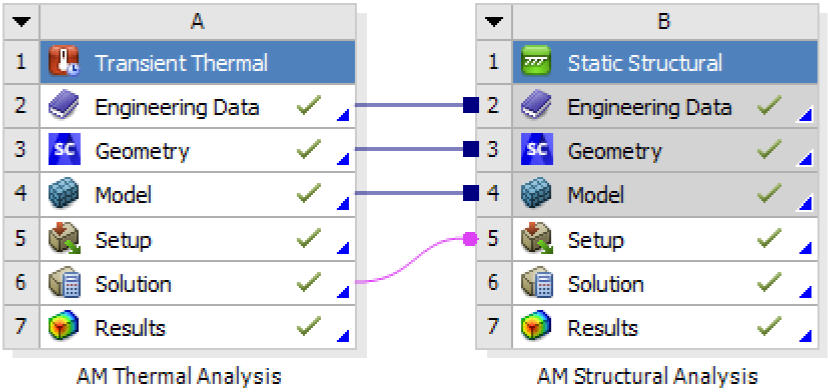

Batch Workbench Job

- If you don't have RSM, you can still submit a whole workbench project as a single job.
  - Archive the project into a wbpz file (optional).
  - Move that file onto cluster
  - Extract using workbench (or just unzip it, wbpz is just a gzip file):

```
/ansys_inc/v202/Framework/bin/Linux64/runwb2 -B -E 'Unarchive(ArchivePath=r"Thermal Structural.wbpz", ProjectPath=r"Thermal Structural.wbpj", Overwrite = True)’
```

  - Update the whole project:

```
/ansys_inc/v202/Framework/bin/Linux64/runwb2 -B -E 'Update(); Save(Overwrite=True)' -F "Thermal Structural.wbpj“
```

  - Learning Workbench scripting language is not covered here, but the “recording” feature is highly recommended.

```
-A addinsfile: Use addins file 'addins'
-B           : Run in batch
-C configfile: Use configuration 'config'
-D dataFile  : Import an application data file.
-E statement : Execute a journal statement at startup.  May be repeated.
-F WBProject : Open a Workbench project file (*.wbpj)
-I           : Run interactively
-L language  : Show UI in 'language'
-R replayfile: Use replay file 'replay'
-W workspace : Start UI with specified workspace
-X           : Run interactively & exit upon journal replay end
```

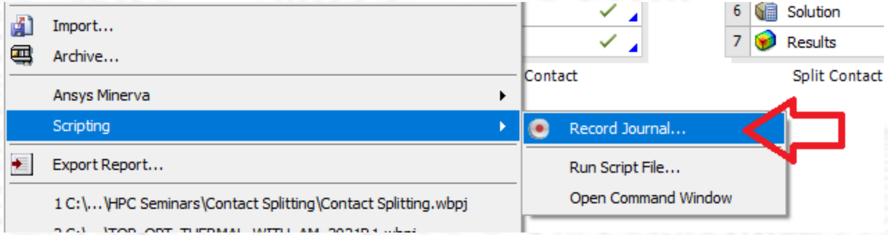

Single Batch Solve

- Alternatively, submit your solves to mechanical one by one:

/ansys_inc/v202/ansys/bin/ansys202 -b -dis -np {nprocs} -i ds.dat -o solve.out

-b	Batch mode
-dis	Distributed memory
-i	APDL input file
-o	Solver output file
-np	Number of processes (cores)
-mpifile {filename}	file containing lists of processes on per host bases for multiple nodes

HPC1
HPC1
HPC2
HPC2
etc.

Interactive Cluster Jobs

- Many job schedulers allow interactive job sessions, including X forwarding (graphics).
- For PBS:

```
qsub -I -X -N Jobname -l select=16
```

- Then just open Workbench and use the GUI:

```
/ansys_inc/v202/Framework/bin/Linux64/runwb2
```

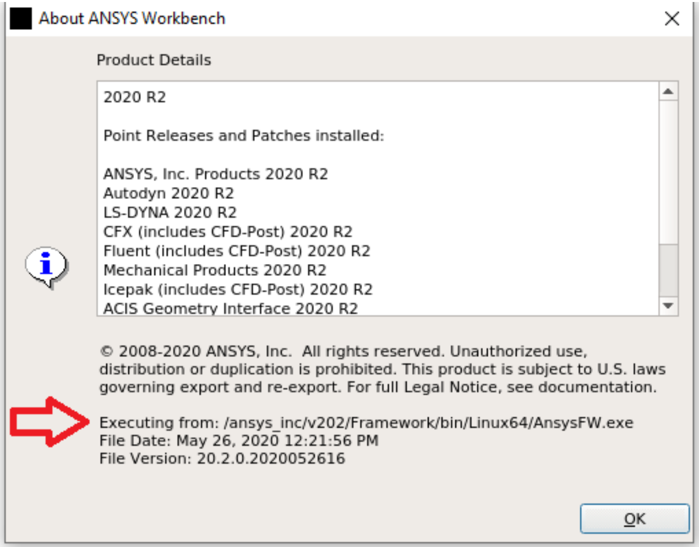

Regarding File Sizes

- Mechanical can create tons of data that isn't always required.
- Partial files, starting with file#.*, can usually be deleted.
- Reduce data being saved in Analysis Settings > Output Controls
  - Particularly save frequency (store results at)
  - Also choose outputs required
- Enable results file compression, or compress results via zip, xz, etc. (60%)
- Clear results in unneeded systems before retrieval:
  - Modal being used before random vibration or response spectrum etc.
  - Transient thermal before structural
- Disable restart points
- Delete unneeded files

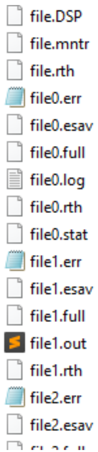

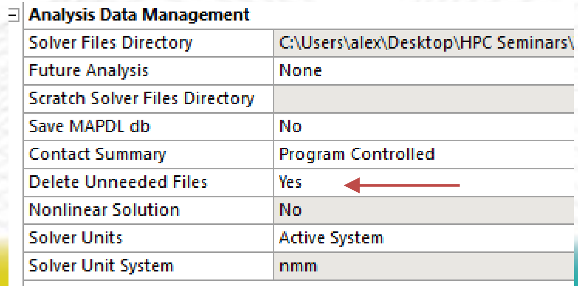

Analysis Data Management	
Solver Files Directory	C:\Users\alex\Desktop\HPC Seminars\
Future Analysis	None
Scratch Solver Files Directory	
Save MAPDL db	No
Contact Summary	Program Controlled
Delete Unneeded Files	Yes
Nonlinear Solution	No
Solver Units	Active System
Solver Unit System	nmm

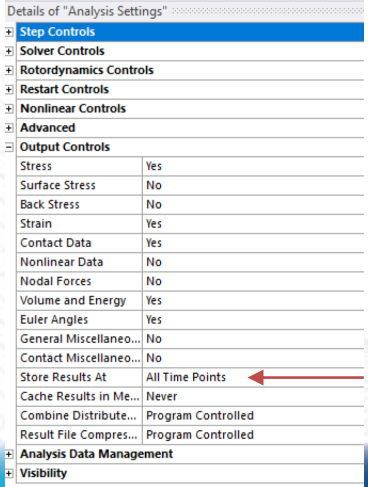

Details of "Analysis Settings"

Step Controls	
Solver Controls	
Rotordynamics Controls	
Restart Controls	
Nonlinear Controls	
Advanced	
Output Controls	
Stress	Yes
Surface Stress	No
Back Stress	No
Strain	Yes
Contact Data	Yes
Nonlinear Data	No
Nodal Forces	No
Volume and Energy	Yes
Euler Angles	Yes
General Miscellaneo...	No
Contact Miscellaneo...	No
Store Results At	All Time Points
Cache Results in Me...	Never
Combine Distribute...	Program Controlled
Result File Compres...	Program Controlled
Analysis Data Management	
Visibility	

Reading Solver Output

- Which HPC settings are being used is right near the start of the solve:

```
         *****  ANSYS COMMAND LINE ARGUMENTS  *****
 BATCH MODE REQUESTED (-b)     = NOLIST
 INPUT FILE COPY MODE (-c)     = COPY
 DISTRIBUTED MEMORY PARALLEL REQUESTED
     32 PARALLEL PROCESSES REQUESTED WITH SINGLE THREAD PER PROCESS
   TOTAL OF    32 CORES REQUESTED
```

- Next look for solution options and you can find solver requested (it can change!):

```
                 S O L U T I O N   O P T I O N S

   PROBLEM DIMENSIONALITY. . . . . . . . . . . . .3-D
   DEGREES OF FREEDOM. . . . . . UX   UY   UZ   ROTX ROTY ROTZ
   ANALYSIS TYPE . . . . . . . . . . . . . . . . .STATIC (STEADY-STATE)
   OFFSET TEMPERATURE FROM ABSOLUTE ZERO . . . . .  273.15
   EQUATION SOLVER OPTION. . . . . . . . . . . . .PCG
      TOLERANCE. . . . . . . . . . . . . . . . . . 1.00000E-08
   GLOBALLY ASSEMBLED MATRIX . . . . . . . . . . .SYMMETRIC
```

```
                 S O L U T I O N   O P T I O N S

   PROBLEM DIMENSIONALITY. . . . . . . . . . . . .3-D
   DEGREES OF FREEDOM. . . . . . UX   UY   UZ
   ANALYSIS TYPE . . . . . . . . . . . . . . . . .STATIC (STEADY-STATE)
   OFFSET TEMPERATURE FROM ABSOLUTE ZERO . . . . .  273.15
   NONLINEAR GEOMETRIC EFFECTS . . . . . . . . . .ON
   EQUATION SOLVER OPTION. . . . . . . . . . . . .SPARSE
   NEWTON-RAPHSON OPTION . . . . . . . . . . . . .PROGRAM CHOSEN
   GLOBALLY ASSEMBLED MATRIX . . . . . . . . . . .SYMMETRIC
```

- Domain decomposer tells us nodes and elements as well as load balance ratio:

```
   D I S T R I B U T E D   D O M A I N   D E C O M P O S E R

 ...Number of elements: 293650
 ...Number of nodes:    298729
 ...Decompose to 32 CPU domains
 ...Element load balance ratio =     8.977
```

Reading Solver Output

- For direct solver there is a memory usage prediction then a confirmation as to if we solving in RAM (in-core) or by paging (out-of-core):

```
DISTRIBUTED SPARSE MATRIX DIRECT SOLVER.
 Number of equations =      894111,    Maximum wavefront =    294

 Local memory allocated for solver                   =    191.089 MB
 Local memory required for in-core solution          =    182.987 MB
 Local memory required for out-of-core solution =     78.563 MB

 Total memory allocated for solver                   =   6291.312 MB
 Total memory required for in-core solution          =   6023.870 MB
 Total memory required for out-of-core solution =   2484.010 MB

*** NOTE ***                                CP =      30.423   TIME= 02:20:54
The Distributed Sparse Matrix Solver is currently running in the
in-core memory mode.  This memory mode uses the most amount of memory
in order to avoid using the hard drive as much as possible, which most
often results in the fastest solution time.  This mode is recommended
if enough physical memory is present to accommodate all of the solver
data.
```

- Important output about convergence and progress:

```
*** LOAD STEP     3   SUBSTEP     5  COMPLETED.    CUM ITER =     14
*** TIME =   2.50000          TIME INC =  0.100000
*** AUTO STEP TIME:  NEXT TIME INC = 0.10000      UNCHANGED
```

```
*****************************************************
*************** FINISHED SOLVE FOR LS 1 *************
*****************************************************
****************** SOLVE FOR LS 2 OF 2 ***************
```

Reading Solver Output

- End of solve summary.
- Pay attention to time spent computing solution when doing quick benchmarks.

```
Total CPU time for main thread                    :      110.5 seconds
Total CPU time summed for all threads             :      132.2 seconds

Elapsed time spent obtaining a license            :        0.6 seconds
Elapsed time spent pre-processing model (/PREP7)  :        0.8 seconds
Elapsed time spent solution - preprocessing       :        8.1 seconds
Elapsed time spent computing solution             :      102.8 seconds
Elapsed time spent solution - postprocessing      :        0.4 seconds
Elapsed time spent post-processing model (/POST1) :        0.0 seconds

Equation solver used                              :            Sparse (symmetric)
Equation solver computational rate                :      549.3 Gflops
Equation solver effective I/O rate                :     1005.6 GB/sec

Maximum total memory used                         :     9492.0 MB
Maximum total memory allocated                    :    34880.0 MB
Total physical memory available                   :        252 GB
Maximum total memory available (all machines)     :        503 GB

Total amount of I/O written to disk               :       22.5 GB
Total amount of I/O read from disk                :       53.8 GB

+------ E N D   D I S T R I B U T E D   A N S Y S   S T A T I S T I C S -------+


 *-----------------------------------------------------------------------------*
 |                                                                             |
 |                       DISTRIBUTED ANSYS RUN COMPLETED                       |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 | Ansys 2020 R2            Build 20.2          UP20200601      LINUX x64      |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 | Database Requested(-db)  1024 MB    Scratch Memory Requested       1024 MB  |
 | Maximum Database Used     350 MB    Maximum Scratch Memory Used     387 MB  |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 |        CP Time      (sec) =        132.159       Time  =  02:22:34          |
 |        Elapsed Time (sec) =        117.000       Date  =  03/01/2021        |
 |                                                                             |
 *-----------------------------------------------------------------------------*
```

Reading Solver Output

- End of solve summary.
- Pay attention to time spent computing solution when doing quick benchmarks.

```
Total CPU time for main thread                    :      110.5 seconds
Total CPU time summed for all threads             :      132.2 seconds

Elapsed time spent obtaining a license            :        0.6 seconds
Elapsed time spent pre-processing model (/PREP7)  :        0.8 seconds
Elapsed time spent solution - preprocessing       :        8.1 seconds
Elapsed time spent computing solution             :      102.8 seconds
Elapsed time spent solution - postprocessing      :        0.4 seconds
Elapsed time spent post-processing model (/POST1) :        0.0 seconds

Equation solver used                              :            Sparse (symmetric)
Equation solver computational rate                :      549.3 Gflops
Equation solver effective I/O rate                :     1005.6 GB/sec

Maximum total memory used                         :     9492.0 MB
Maximum total memory allocated                    :    34880.0 MB
Total physical memory available                   :        252 GB
Maximum total memory available (all machines)     :        503 GB

Total amount of I/O written to disk               :       22.5 GB
Total amount of I/O read from disk                :       53.8 GB

+------ E N D   D I S T R I B U T E D   A N S Y S   S T A T I S T I C S -------+


 *-----------------------------------------------------------------------------*
 |                                                                             |
 |                        DISTRIBUTED ANSYS RUN COMPLETED                      |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 | Ansys 2020 R2            Build 20.2          UP20200601      LINUX x64      |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 | Database Requested(-db)  1024 MB    Scratch Memory Requested       1024 MB  |
 | Maximum Database Used     350 MB    Maximum Scratch Memory Used     387 MB  |
 |                                                                             |
 |-----------------------------------------------------------------------------|
 |                                                                             |
 |        CP Time      (sec) =        132.159       Time  =  02:22:34          |
 |        Elapsed Time (sec) =        117.000       Date  =  03/01/2021          |
 |                                                                             |
 *-----------------------------------------------------------------------------*
```

Reading Solver Output

- Enable performance details with:

  DSPOPTION,,,,,,PERFORMANCE

```
DISTRIBUTED PCG SOLVER SOLUTION STATISTICS

  NUMBER OF ITERATIONS=          275
  NUMBER OF EQUATIONS =        24372
  LEVEL OF CONVERGENCE=            1
  CALCULATED NORM     = 0.68121E-08
  SPECIFIED TOLERANCE = 0.10000E-07
  TOTAL CPU TIME (sec)=         0.47
  TOTAL WALL TIME(sec)=         0.57
  TOTAL MEMORY (GB)   =         0.01

   EQUIL ITER   1   CPU TIME =  0.8750      ELAPSED TIME =  0.7277

NUMBER OF SMP THREADS USED   =  1
ELEMENT RESULTS CP TIME      =     0.047
ELEMENT RESULTS ELAPSED TIME =     0.043

  *** ELEMENT RESULT CALCULATION TIMES
    TYPE    NUMBER   ENAME      TOTAL CP   AVE CP

       1      1690  SOLID186      0.219   0.000129
       2       169  SURF154       0.000   0.000000
       3       130  CONTA174      0.000   0.000000

  *** NODAL LOAD CALCULATION TIMES
    TYPE    NUMBER   ENAME      TOTAL CP   AVE CP

       1      1690  SOLID186      0.031   0.000018
       2       169  SURF154       0.000   0.000000
       3       130  CONTA174      0.000   0.000000
*** LOAD STEP     1   SUBSTEP     1  COMPLETED.    CUM ITER =        1
*** TIME =   1.00000         TIME INC =   1.00000      NEW TRIANG MATRIX
```

```
adjust near zero pivots tolerance =   4.356441989940202E-012
# of positive pivots    =        894457
# of negative pivots    =             0
# of near zero pivots   =             0
Maximum pivot           =    0.1929E+07
Minimum pivot           =    0.2311E+03

# of 1x1 pivots         =             0
# of 2x2 pivots         =             0
# of delayed pivots     =             0
# of defered pivots     =             0
Distributed sparse solver maximum pivot= 1928754.37 at node 296908 UZ.
Distributed sparse solver minimum pivot= 231.116041 at node 162699 UZ.
Distributed sparse solver minimum pivot in absolute value= 231.116041
at node 162699 UZ.

===========================
= multifrontal statistics =
===========================

    number of equations                         =          894457
    no. of nonzeroes in lower triangle of a     =        33498628
    no. of nonzeroes in the factor l            =       491252270
    ratio of nonzeroes in factor (min/max)      =          0.6795
    number of super nodes                       =           31183
    maximum order of a front matrix             =            3648
    maximum size of a front matrix              =         6655776
    maximum size of a front trapezoid           =         3576855
    no. of floating point ops for factor        =      6.9664D+11
    no. of floating point ops for solve         =      9.4507D+09
    ratio of flops for factor (min/max)         =          0.5557
    near zero pivot monitoring activated
    number of pivots adjusted                   =               0
    negative pivot monitoring activated
    number of negative pivots encountered       =               0
    factorization panel size                    =             128
    number of cores used                        =              32
    time (cpu & wall) for structure input       =        0.110000        0.112768
    time (cpu & wall) for ordering              =        1.390000        1.426438
    time (cpu & wall) for other matrix prep     =        0.140000        0.143275
    time (cpu & wall) for value input           =        0.030000        0.034700
    time (cpu & wall) for matrix distrib.       =        0.180000        0.168110
    time (cpu & wall) for numeric factor        =        1.300000        1.312331
    computational rate (gflops) for factor      =      535.874410      530.839350
    time (cpu & wall) for numeric solve         =        0.060000        0.057471
    computational rate (gflops) for solve       =      157.511490      164.443687
    effective I/O rate (GB/sec) for solve       =      586.822603      612.649098

    i/o stats: unit-Core          file length              amount transferred
                                words        mbytes          words        mbytes
                     ----     ----------    --------     ----------    --------
                   95-   0       425984.       3. MB        401266.       3. MB
                   95-   1       360448.       3. MB        357662.       3. MB
                   95-   2       360448.       3. MB        359622.       3. MB
                   95-   3       393216.       3. MB        361609.       3. MB
```

Beta Features

- Check out beta features meant to improve efficiency

v202

Beta feature selection

- ☐ Load Case Manager
- ☐ Allow Body Grouping Control for meshes from External Model
- ☐ Skip processing plies from External Model
- ☐ Allow Element Scoping for Material Assignment
- ☐ Allow thermal variation along shell thickness
- ☐ Remote Post Processing
- ☐ Topology Optimization - Shape Optimization
- ☐ Recording
- ☐ Add CAD Files and Delete Parts
- ☐ Select Files To Download On RSM Solve
- ☐ Rigid Dynamics - Expose advanced settings
- ☐ Allow Coupled Field Harmonic and Modal for Piezo Electric
- ☐ Topology Optimization - Enable Back To CAD export of Model.
- ☐ Enable advanced reduction method for user defined criterion
- ☐ Use cable higher order element
- ☐ Enable DPF evaluation
- ☐ Bottom Up Substructuring
- ☐ Trace Mapping Support for LS-DYNA
- ☐ Allow Usage of LS-DYNA Materials through External Model

v211

Beta feature selection

- ☐ Load Case Manager
- ☐ External Model - Allow Body Grouping Control for meshes
- ☐ External Model - Skip processing plies
- ☐ External Model - Send surface loads using SFE
- ☐ External Model - Allow Usage of LS-DYNA Materials
- ☐ Allow Element Scoping for Material Assignment
- ☐ Allow thermal variation along shell thickness
- ☐ Remote Post Processing
- ☐ Topology Optimization - Shape Optimization
- ☐ Record Unpublished API
- ☐ Add CAD Files and Delete Parts
- ☐ Rigid Dynamics - Expose advanced settings
- ☐ Rigid Dynamics - AQWA/RBD/AERODYN cosimulation
- ☐ Topology Optimization - Enable Back To CAD export of Model.
- ☐ Topology Optimization - Enable Experimental features
- ☐ Enable advanced reduction method for user defined criterion
- ☐ Enable DPF evaluation
- ☐ Penalty Based Contact Between SPH Bodies and MLS Formulation for LS-DYNA
- ☐ EXD SPH - Additional Features
- ☐ Allow Acoustics in Coupled Field Transient
- ☐ Create Script Defined Results
- ☐ Create Spot Weld groups
- ☐ EXD Automatic Dynamic Relaxation
- ☐ Structured ALE Support for LS-DYNA
- ☐ Enable LS-DYNA Implicit Features
- ☐ Enable experimental fast animation for basic results (Windows Only)

Explicit Sovers

- Very similar in behaviour to iterative solvers
  - Low RAM usage
  - No AVX requirements (as shipped)
  - Scale very well (better actually)

Contact Efficiency

- Demos of contact splitting and trimming

Questions?

Thanks for listening!

Apickard@simutechgroup.ca

www.SimuTechGroup.com/Offices