ANSYS HPC Seminar Series

Electronics Desktop

Prepared and presented by

Alex Pickard (HPC) apickard@simutechgroup.ca
Zoe Zhu (HF) zzhu@simutechgroup.ca
Natalya Vorobtsova (LF) nvorobtsova@simutechgroup.com

Who Are We

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

Today’s Adenda

- Introduction to HPC, Electronics Desktop HPC, solvers
- Model Decomposition Techniques (parallelization methods)
- Sizing resource request for your model
- Intro to clusters and Job Schedulers
- Batch solving on clusters
- Other cluster job submission techniques
- Simulation best practices for performance and scalability
- Reading progress from solver logs
- Academic licence changes
- Computer platform recommendations (CPU models & features, RAM quantity, memory bandwidth, storage, networking)
- Hardware recommendations for users (solver and preprocessor)

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.

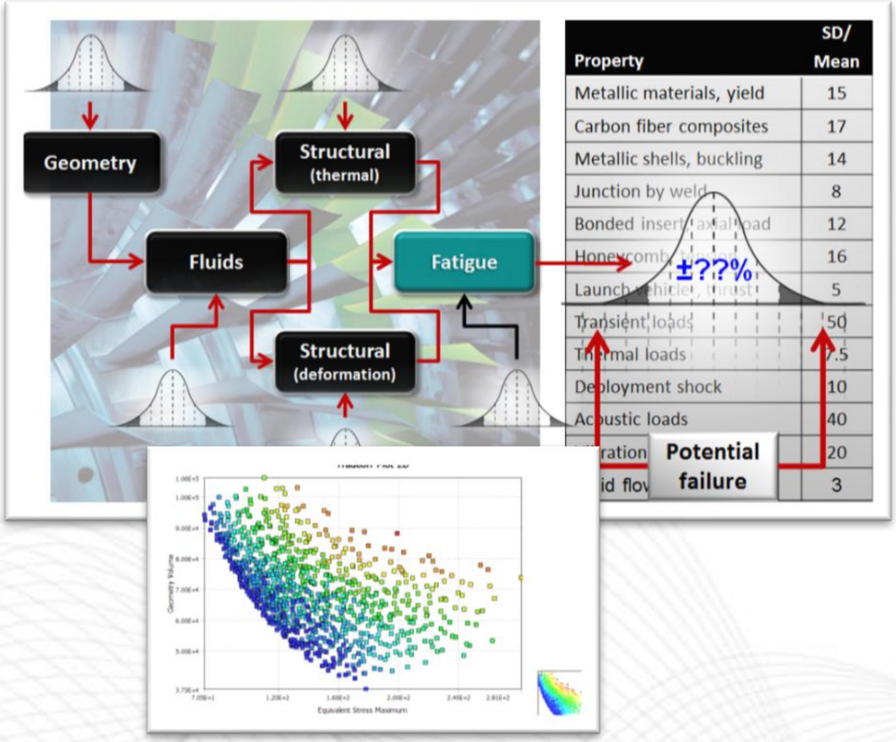

What is HPC?

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.

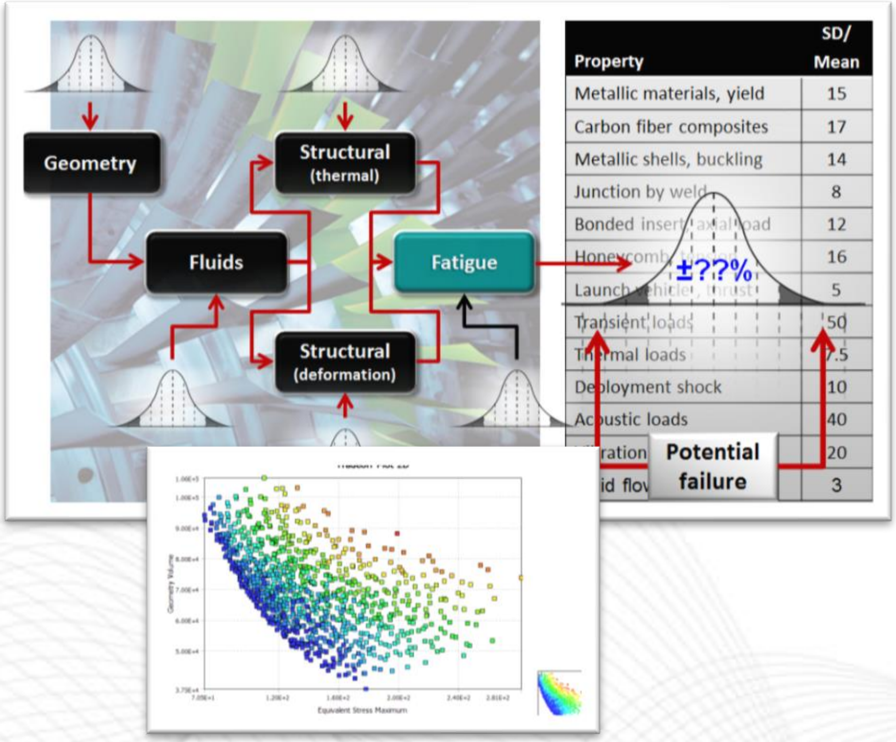

ANSYS Electronics Desktop

- The ANSYS Electronics Desktop is a suite of solvers built into a single analysis tool that can be used for a wide range of low frequency and high frequency electronics problems.
  - Antennas
  - Integrated Circuits
  - Electrostatics
  - Magnetostatics
  - Eddy current
  - Electronics Cooling
  - Radio Frequency Interference

ANSYS EDT Solvers

- The solvers included in Electronics Desktop are:
  - **HFSS:** High frequency solver for antennas, Integrated Circuits,
  - **Maxwell:** Low Frequency Solver
  - Q3D: Quasistatic 3D solver
  - Circuit: Schematic based circuit simulator for RF/SI
  - RMxprt: Template based machine design tool
  - **Icepak:** Thermal performance of electronics
  - EMIT Radio Frequency Interference Prediction
  - Twin Builder Multi-domain mixed-signal simulator

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:
  - Remeshing for example

Example of Spatial Decomposition:

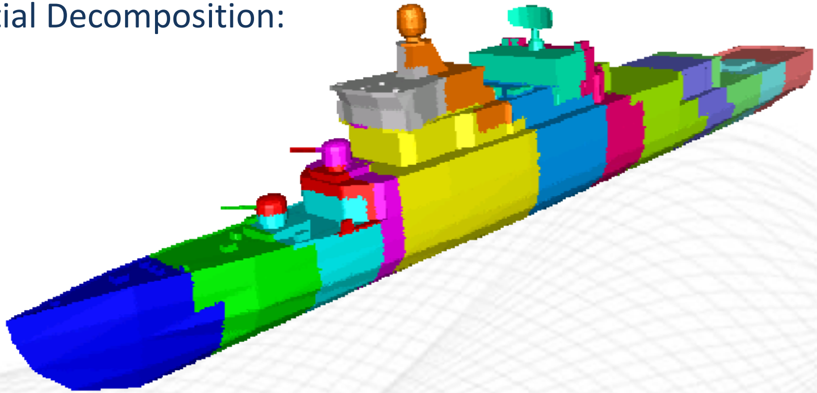

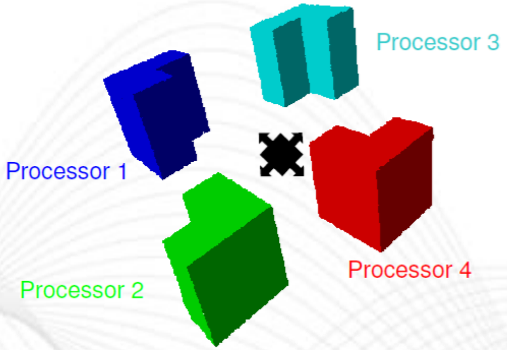

Maxwell HPC Methods

- Maxwell HPC supports all steps being parallelized:
  - Meshing, matrix assembly, solving, post processing
- All Maxwell solvers are supported.
- In its classic matrix decomposition mode early performance saturation is achieved (around 8 cores)
- Other parallelization techniques are needed to utilize more resources efficiently:
  - Time Domain Decomposition
  - Distributed Solve Options (parameterization)

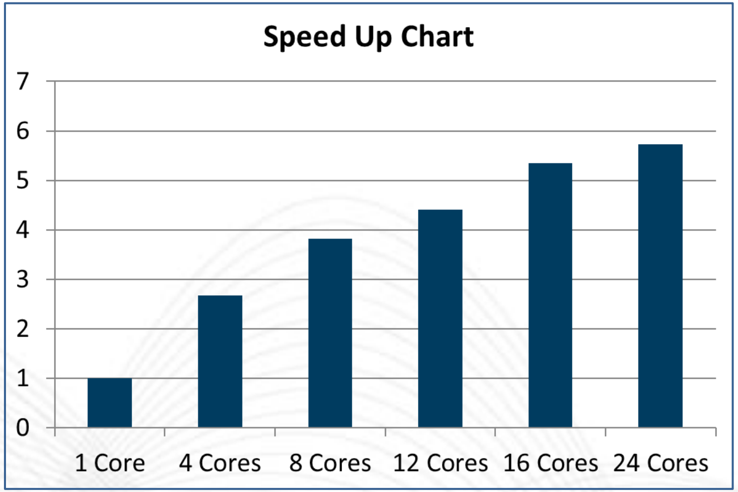

Time Domain Decomposition Method (TDM)

- For transient simulations, time domain decomposition allows solving of all time steps simultaneously instead of sequentially.

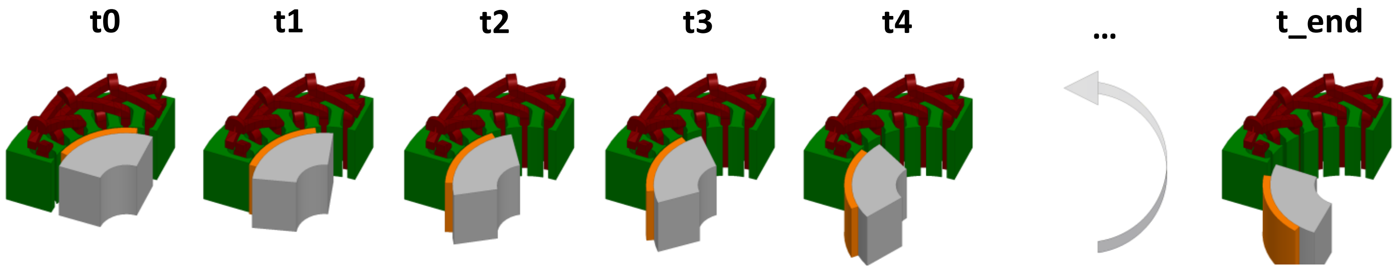

- Very loose coupling in time so timesteps are relatively independent.
- Solving multiple time points simultaneously allows greater parallel efficiency.

Time Domain Decomposition Method (TDM)

- Spatial divisions are created and then all timesteps are added that domains matrix.
- Subdivisions are solved in series, but the time steps are distributed over all tasks.
- Use minimum subdivisions (largest chunks), up to the limits of RAM.

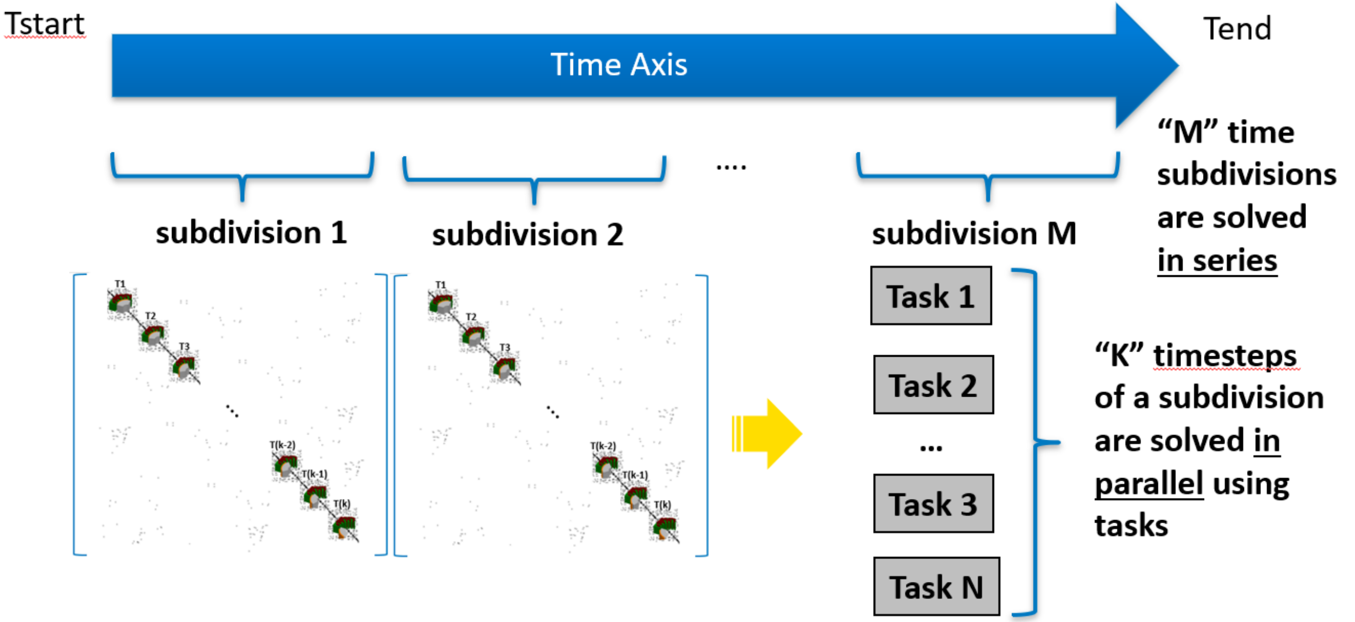

Distributed Solve Option (DSO/LSDSO)

- DSO allows multiple parametric variations of one design to be solved independently.
  - No interdependence so ideal scaling potential.
- High level and low level communication
  - Less network performance dependent

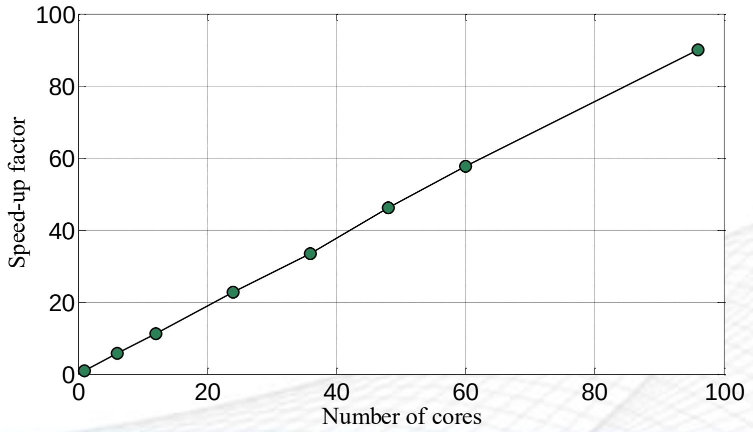

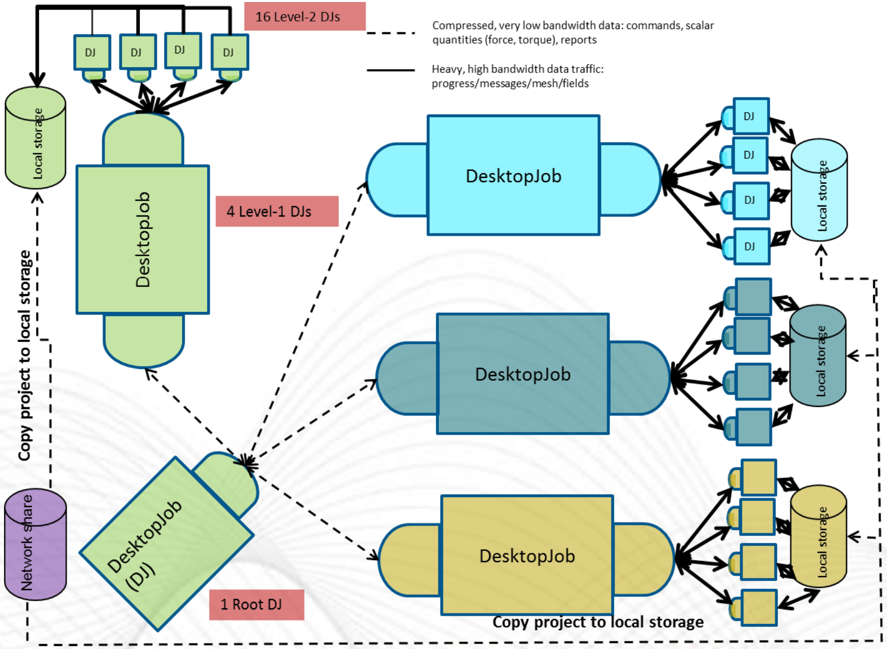

HFSS HPC Methods

- HFSS supports multiple steps being parallelized:
  - Meshing, matrix assembly, solving, field recovery
- The classic **Matrix Multiprocessing** technique gives each core its own frontal matrix.
  - Scales to ~10-20 cores
- **Domain Decomposition Method** (DDM) breaks FE mesh into domains and solves as distributed solver.
- **Periodic Domain Decomposition** is DDM applied to periodic structures like antenna to virtually duplicate geometry.
  - Significant speedup from reusing mesh.
- **Spectral Decomposition Method** (SDM) accelerates frequency sweeps by distributing multiple frequency tasks in parallel.
  - Can be combined with others.
  - Mixed networking requirements due to "independent" tasks
- **Distributed Matrix Solver** uses distributed memory techniques on an Integral Equation problem where the matrix is distributed, rather than the FE mesh.
- **Hybrid Domain Decomposition Method** allows mixed models of Finite Elements and Integral Equations to be solved separately, in parallel.
  - Above techniques can be used on each subdomain.
- **Design Space Optimization** (parameters) allows further parallelization.

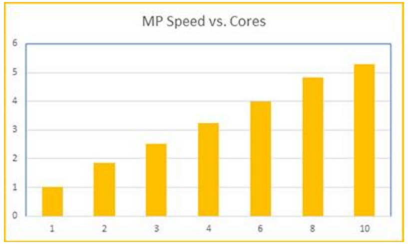

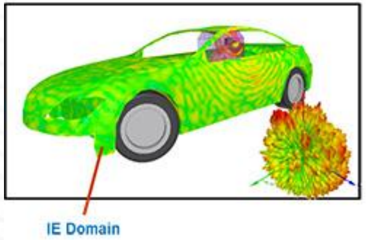

Solver Memory Requirements

- The multitude of parallelization techniques available all require increased memory usage to increase performance.
  - Solving more tasks (frequencies, timesteps, DSO, etc.) required RAM for each task.
- The use of adaptive meshing techniques makes predicting memory requirements incredibly difficult, but once the adaptive passes are done, the memory used is reported.
  - Some variance is required for each subsequent task.
- To give some foresight into resource requirements, do adaptive passes and look in solver output for memory usage after convergence, then plan from there (or use auto).
- If enough RAM is available, best results are approximately 4 cores / task, up to machine limit.

```
Adaptive Pass 10
Adaptive Passes converged
Simulation Summary:
…
Adaptive Meshing : Elapsed time: 00:02:12, total memory: 8.37 GB
max solved tets: 109152, max matrix size: 837932, max bandwidth:  51.4
```

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:
  - PBS Pro (Linux)
  - Torque (Linux, very similar to PBS, just had official support dropped but still works)
  - Platform LSF (Linux)
  - Sun Grid Engine (Linux)
  - Windows HPC Server (Windows Server)
- Batch solving can be custom scripted to work with basically any job scheduler.

RSM Job Submission

- Submitting project to RSM via GUI is easier than batch submission.
  - Control over which subsystems are solved, and resource usage
- Project must be stored in a shared location, accessible to both submission workstation and cluster.
  - Relative path translation is available.
- Preview button handy for guiding batch solves:

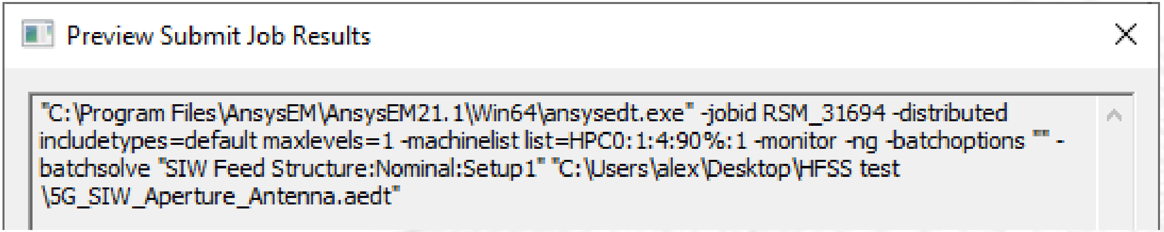

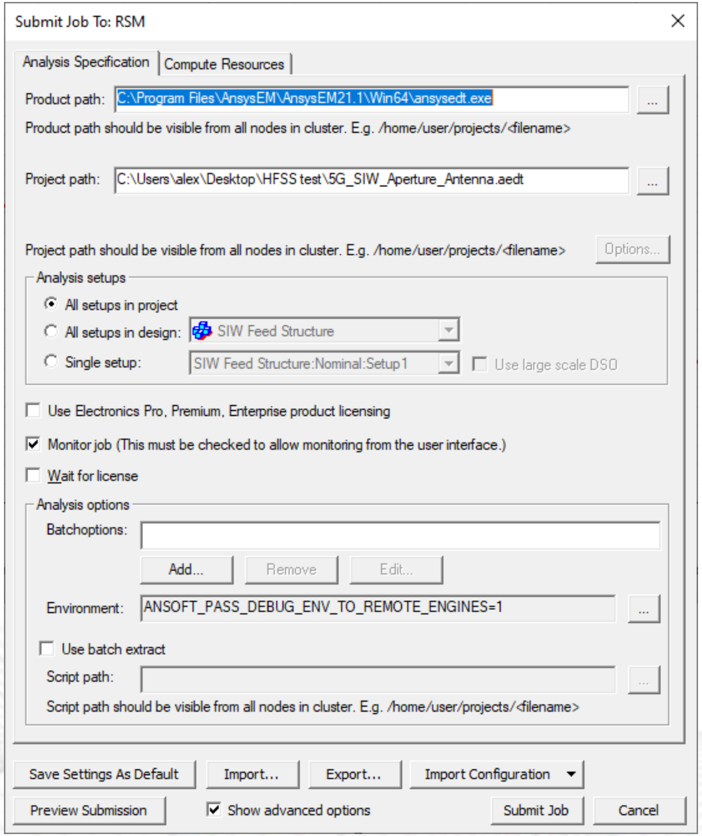

Batch Job

- The batch submission process is more complex than other ANSYS Products.
- The preview button is very helpful if you are unfamiliar with required arguments
- The amount of batchoptions available is staggering, but documented in help.

Some examples:

- ```
  ansysedt -batchoptions "'HFSS/CreateStartingMesh'=1" -batchsolve "D:\projects\MyProject.aedt"
  ```
- ```
  ansysedt -distributed -machinelist list="255.255.1.1,255.255.1.2" -batchsolve
  myDesign:Optimetrics "C:\myProject.aedt"
  ```
- ```
  ansysedt -BatchSolve -Distributed
  -Machinelist file=/home/jsmith/hosts/list2
  -batchoptions "HFSS-IE/Preferences/MemLimitHard=8388608
  HFSS-IE/Preferences/MemLimitSoft=6291456
  HFSS-IE/Preferences/NumberOfProcessors=4
  HFSS-IE/Preferences/NumberOfProcessorsDistributed=1"
  /home/jsmith/projects/project2.aedt
  ```
  ```

Multi-Step Job Submission

- The amount of resources required in cores and RAM ramps up significantly from the start of the solve to the final distributed solution.
  - Minimal cores and RAM for initial mesh generation
  - Modest cores and significant RAM needed for mesh adaptive passes. Sometimes a whole single computer.
  - A whole cluster can be used for the final tasks level parallelization.
- Multi-step submission, either via RSM or batch, allows resource use to scale with this ramp-up.

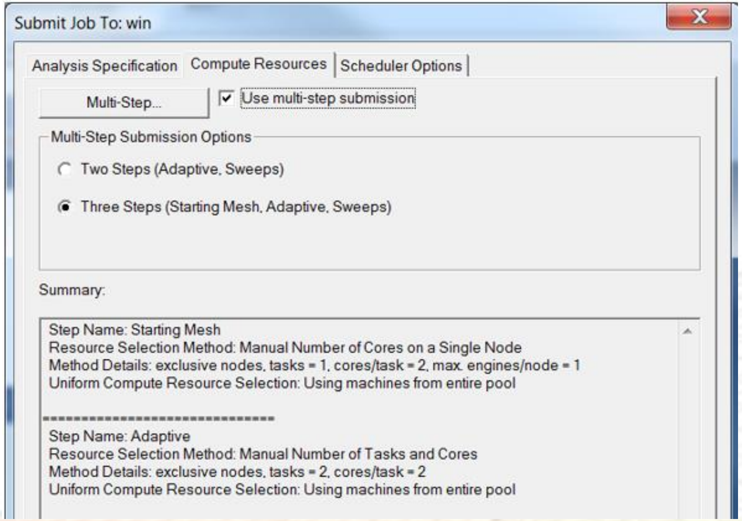

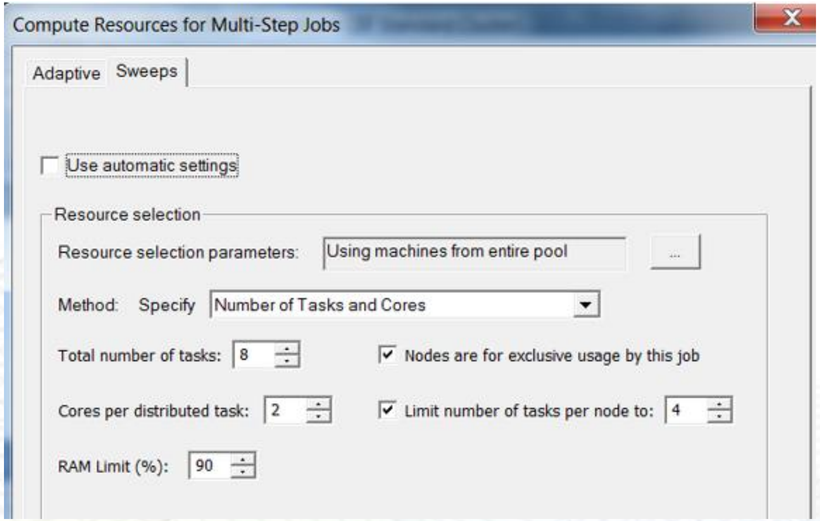

Multi-step Job Submission

- 77GHz Automotive Radar with Package and Radome
  - Simulation specifications
    - Medium sized problem
    - Number of excitations: 8
    - Interpolating Frequency Sweep 401 points.
    - Solution Frequency 77 GHz (Save fields)
    - Total tetrahedra: 238k
    - Matrix size: 15.6M
  - AEU savings
    - Up to 16 % in AEUs savings with the 3 step multi-step submission

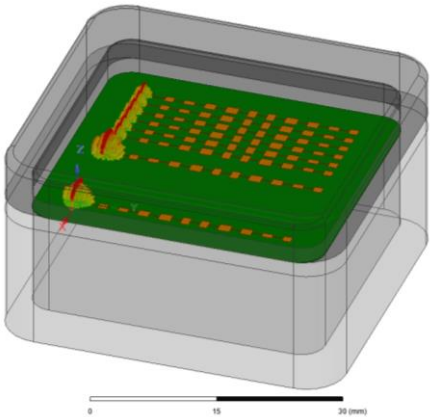

*AECs usage includes Hardware and Software Cost

Ansys Electronics Cloud 2021R1							
Settings	**Confs.**	**Cores**	**RAM [TB]**	**Total Time**	**Speed Up [%]**	**AECs Usage***	**AECs Saving [%]**
Single	L	176	1.4	02:25:09	0	116.80	0
2 step	M / L	88 / 176	.704 / 1.4	02:51:56	-18	98.80	15
3 step	S / M / L	44 / 88 / 176	.352 / .704 / 1.4	03:07:26	-29	98.10	16

Note for Settings:
2 step: Medium configuration was used for starting mesh and adaptive, large configuration was used for frequency sweep.
3 step: Small configuration was used for starting mesh, medium configuration for the adaptive, large configuration was used for frequency sweep.

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 run ansysedt for the GUI:

```
/opt/AnsysEM/AnsysEM21.1/Linux64/ansysedt
```

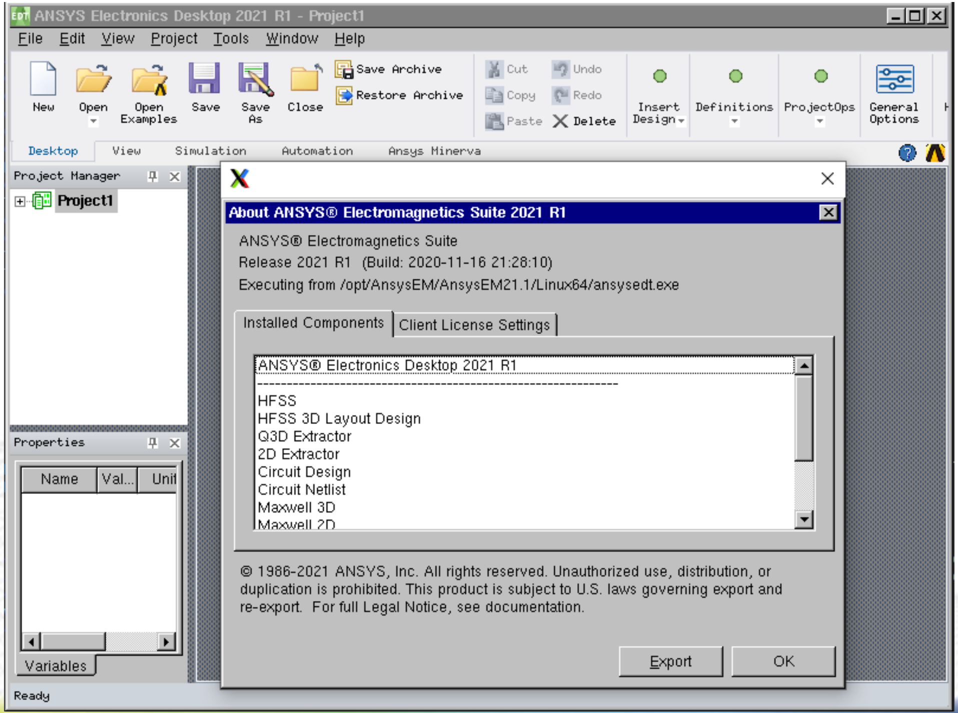

HFSS Job Distribution

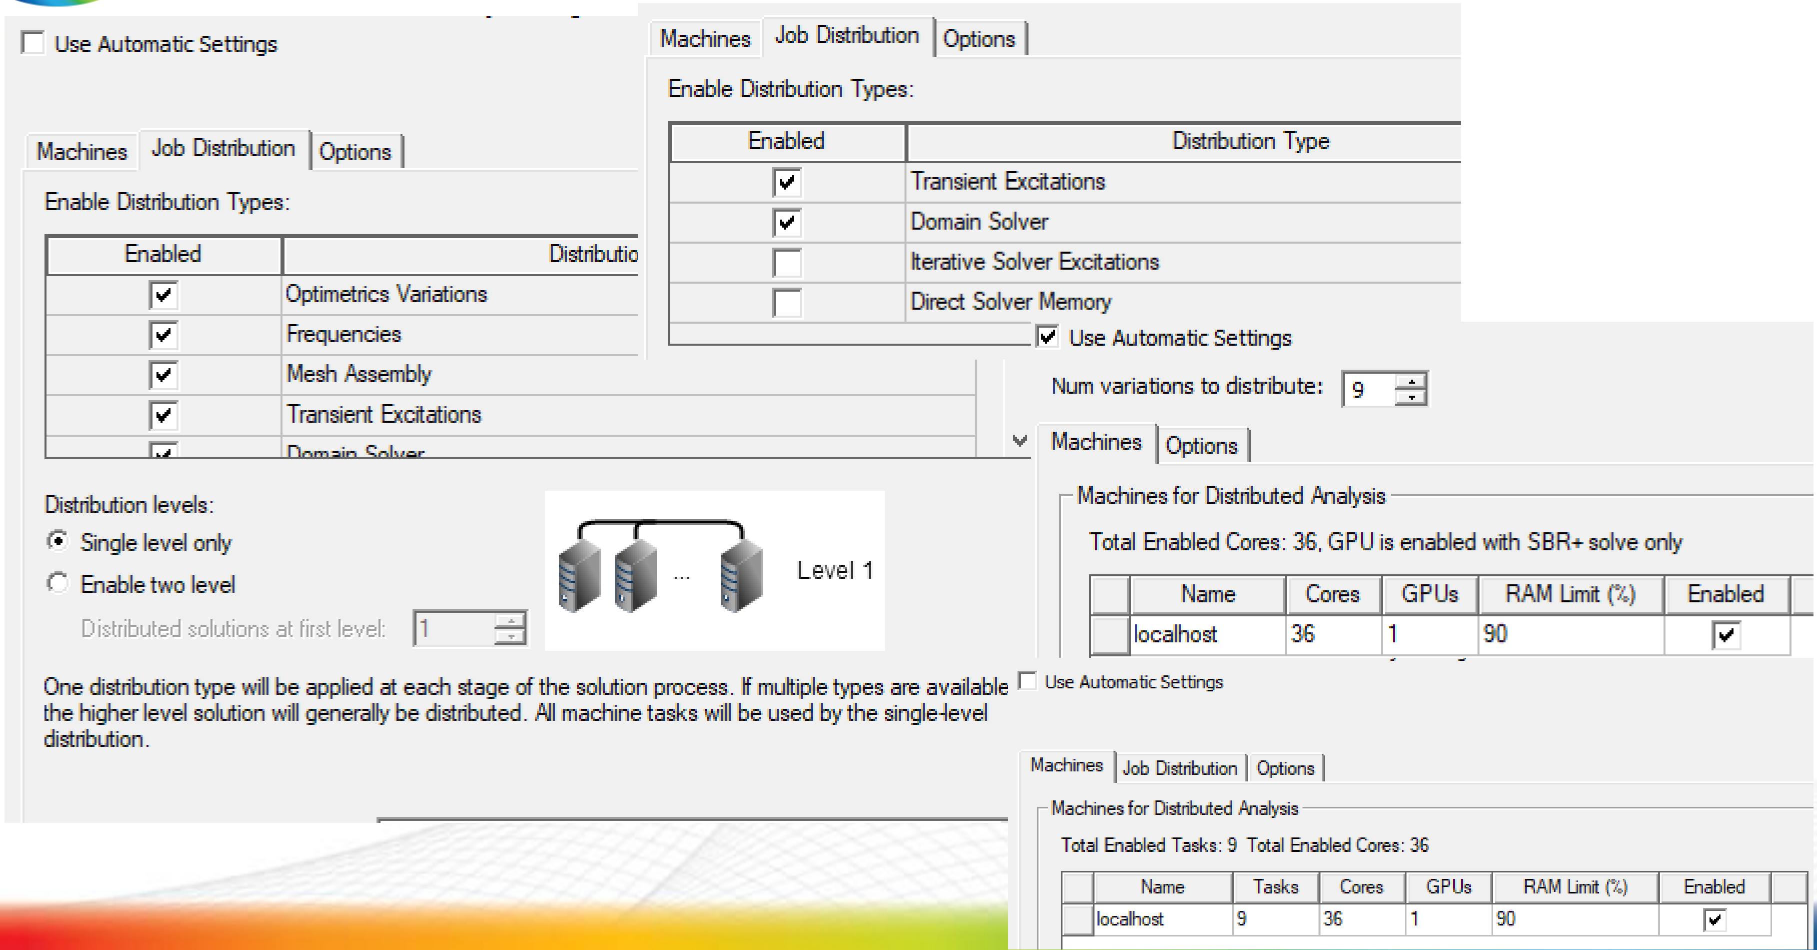

HFSS Simulation Profile

Simulation: Setup1

Design Variation: 3='0.289mm' patchX='4.8mm' patchY='3.52mm' phi_scan='0deg' subH='3.5mil' theta_scan='0deg'

Profile | Convergence | Matrix Data | Mesh Statistics

Task	Real Time	CPU Time	Memory	Information
Adaptive Pass 23				Frequency: 24.125GHz
Mesh (volume, adaptive)	00:00:04	00:00:04	128 M	117124 tetrahedra
Adaptive Meshing Frequency: 24.125GHz on ZoeDesktop.simute...				
Simulation Setup	00:00:02	00:00:02	214 M	Disk = 0 Bytes, 104720 tetrahedra
Matrix Assembly	00:00:05	00:00:14	1.08 G	Disk = 75 Bytes, 104720 tetrahedra , 1 lumped port(s)
Solver DCS6	00:00:13	00:00:43	2.95 G	Disk = 4 Bytes, matrix size 642299 , matrix bandwidth 20.7
Field Recovery	00:00:01	00:00:06	2.95 G	Disk = 3.15 MBytes, 1 excitations , Average Order 0.831761
Data Transfer	00:00:00	00:00:00	76.9 M	Adaptive Pass 23
				Adaptive Passes converged
Adaptive Meshing				Elapsed time: 00:04:15
Frequency Sweep				Time: 12/22/2020 18:19:35
Solution: Sweep				Discrete HFSS Frequency Sweep, Solving Distributed - up to 2 frequencies in parallel
				From 24.0625GHz to 24.1875GHz, 3 Frequencies
				Frequency: 24.125GHz has already been solved
				HPC Enabled

HFSS Mesh Information

Simulation: Setup1

Design Variation: 3='0.289mm' patchX='4.8mm' patchY='3.52mm' phi_scan='0deg' subH='3.5mil' theta_scan='0deg'

Profile | Convergence | Matrix Data | Mesh Statistics

Number of Passes
Completed 23
Maximum 30
Minimum 1

Max Mag. Delta S
Target 0.005
Current 0.0029709

View: (•) Table ( ) Plot

Export...

CONVERGED

Consecutive Passes
Target 1
Current 1

Default Settings
Save Defaults | Clear Defaults

Pass Number	Solved Elements	Max Mag. Delta S
1	13544	N/A
2	16171	0.17653
3	17473	0.20134
4	19496	0.22091
5	20834	0.24897
6	20838	0.11119
7	21850	0.06242
8	24190	0.049296
9	27234	0.03638
10	28203	0.023591
11	29355	0.014164
12	32416	0.014241
13	34631	0.014693
14	39321	0.019396
15	44134	0.006288
16	50540	0.010413
17	53349	0.0094318
18	61429	0.017404
19	69010	0.016957
20	78770	0.0078946
21	86033	0.0084678
22	94761	0.0082414
23	104720	0.0029709

Simulation: Setup1

Design Variation: 3='0.289mm' patchX='4.8mm' patchY='3.52mm' phi_scan='0deg' subH='3.5mil' theta_scan='0deg'

Profile | Convergence | Matrix Data | Mesh Statistics

Total number of elements: 117124

	Num T...	Min edge l...	Max edge l...	RMS edge l...	Min tet v...	Max tet v...	Mean tet ...	Std Dev
sub	67736	0.0151141...	1.54878...	0.138615...	2.0705e...	0.019221...	0.000101...	0.00048
RadiatingSurface	34937	0.0201149...	3.47459...	0.553164...	3.94917...	0.944982...	0.017796...	0.06522
Line1	12404	0.0288975...	0.411846...	0.198861...	8.71327...	0.000379...	4.95541e...	4.53858
PML_RadiatingSurface_1	880	0.752535...	4.56382...	2.05676...	0.02551...	5.82802...	0.433714...	0.60931
PML_RadiatingSurface_5	470	0.759923...	4.29557...	2.25037...	0.00387...	4.21307...	0.536606...	0.61948
PML_RadiatingSurface_3	452	0.639108...	5.39774...	2.2339	0.00642...	8.19784...	0.557976...	0.80305
PML_RadiatingSurface_9	126	1.35015...	6.223	2.92671...	0.10796...	7.59961...	1.21416...	1.48843
PML_RadiatingSurface_7	119	1.46595...	6.223	3.00034...	0.15759...	6.80063...	1.28559...	1.35407

Antenna Array Optimization

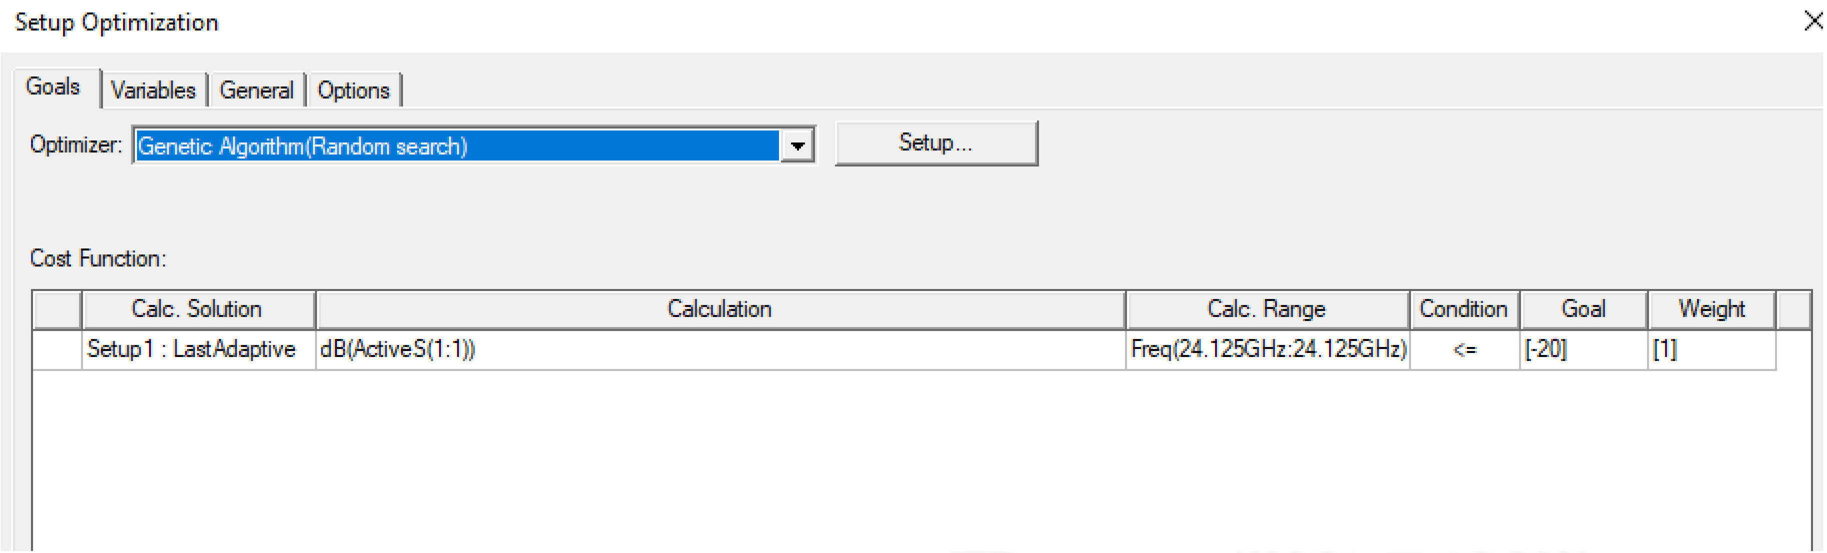

Regions and Boundaries inside HFSS

- The boundary can be determined based on the applications:

Antenna/EMI	Array/RCS	SI
Radiation		**Radiation**
PML	**PML**	
FE-BI	**FE-BI**	

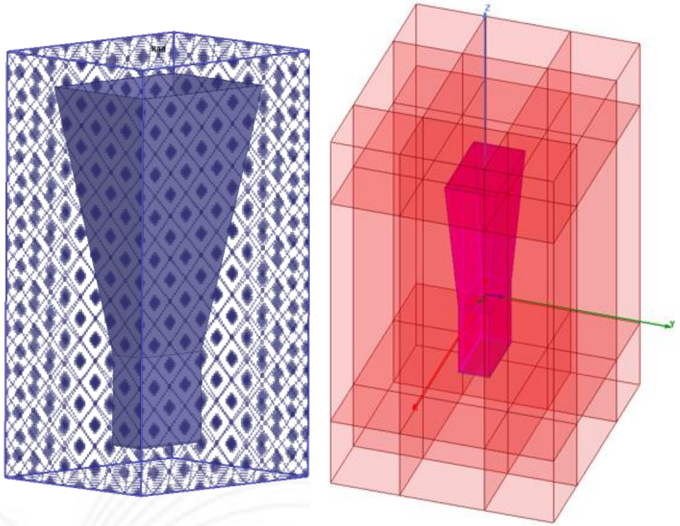

Boundary is $\lambda/4$ away from horn aperture in all directions

- The padding (PD) for the region is determined by the following rules:
  - PD = λ / 3, for radiation boundaries (or ABC)
  - PD = λ / 4, for PML boundaries
  - PD = λ / 8, for FEBI boundaries

Solution Options inside HFSS

- Direct solver and domain decomposition (DDM)
  - Direct solver is optimal for IC/package/board simulations
  - DDM is optimal for antenna/RCS applications
    - This picture shows the mesh needed for the entire jet simulation being split into four pieces that are spread among different computers and connected through DDM
    - Finite array DDM meshes the unit cell using adaptive mesh and duplicates this mesh to all other cells

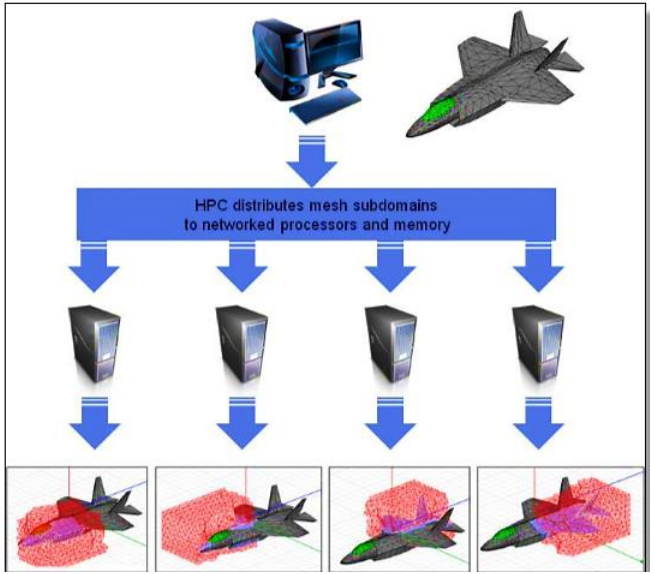

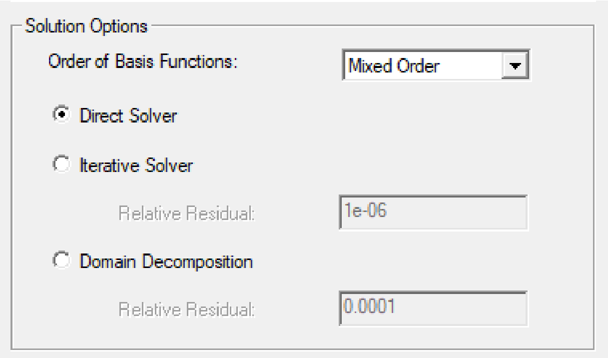

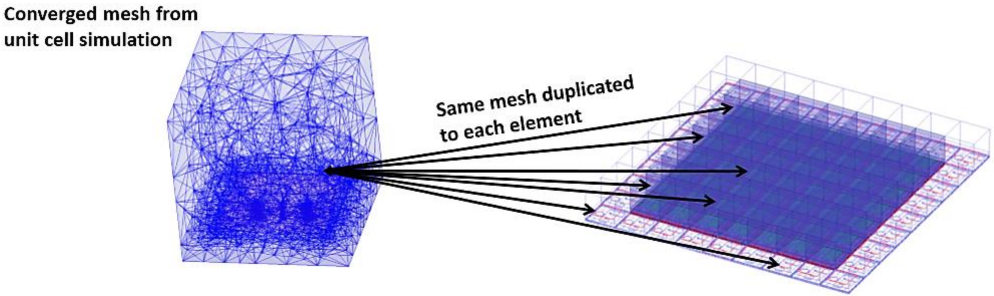

Solution Management

- ***Save radiated fields only*** option enables running large array simulations without filling up the hard disk
- In 8x8 vivaldi array the disk space comparison for with and without ***Save radiated fields only*** option is presented here:
  - 8x8 array -> 64 unit cells
  - 4 excitations per unit cell -> 256 total excitations

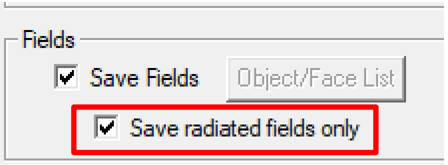

Setup	Disk Space Used	Disk Space Savings
Save all fields (default)	589GB	reference
Save radiated fields only	2.98GB	198x

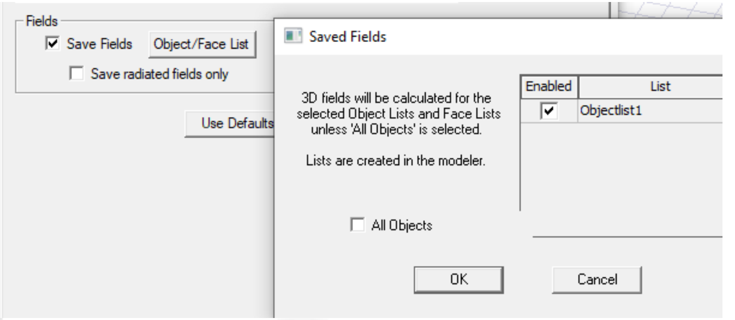

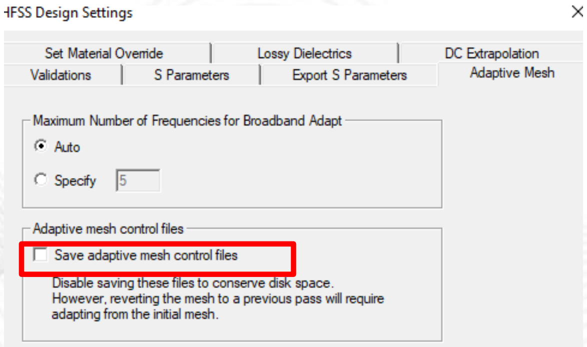

- Save fields on the selected objects only
- Option to not save *.adp files

ANSYS HFSS 2021 R1: Mesh Fusion

- Mesh Fusion **Advantages**
  - **Robust:** Higher overall Mesh Quality
    - Component level mesh settings
    - Mesh tolerance at scale of component
  - **Faster:** Concurrent (i.e. parallel) mesh generation
  - **Scalable:** Mesh larger and more complex complete “Electromagnetic Systems”
  - **No Limits!**
- A Major Breakthrough in HFSS Technology
  - Uncompromised and accurate: Fully coupled fields across region interfaces!
  - Solver delivers the true HFSS *Gold-standard Accuracy*

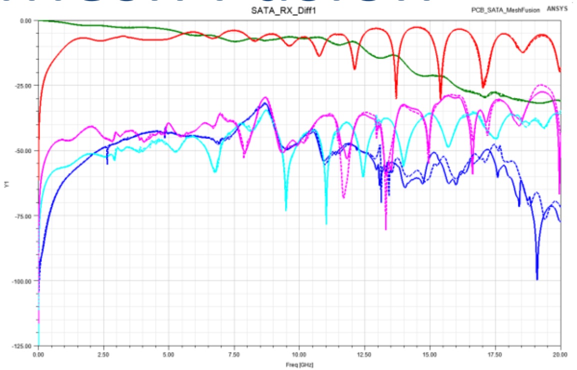

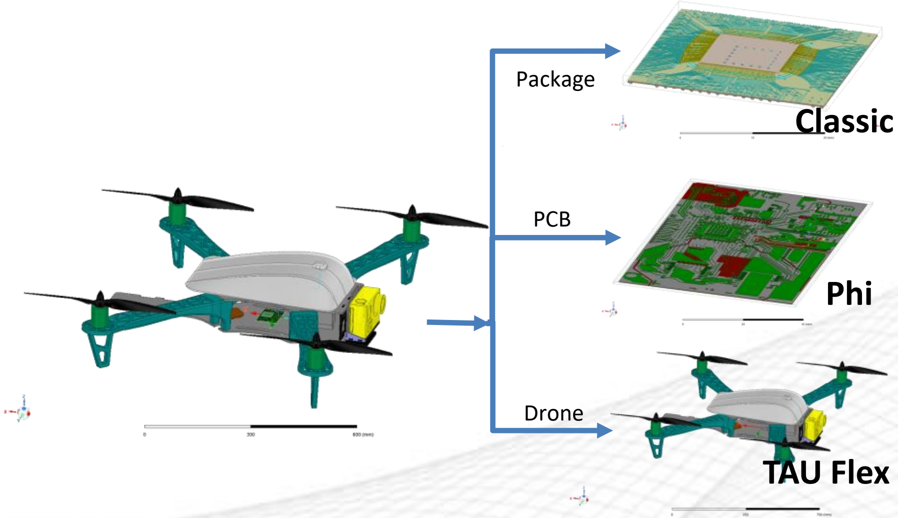

Mesh Assembly Manager

Concurrent Mesh Generation
Create Domain **Contacts**

PCB Volume Mesh

Domain 1
Initial Mesh: Phi

Package Volume Mesh

Domain 2
Initial Mesh: Classic

‘Native’ Volume Mesh

Domain 3
Initial Mesh: Tau Flex

Assembly Volume Mesh

HFSS Solver

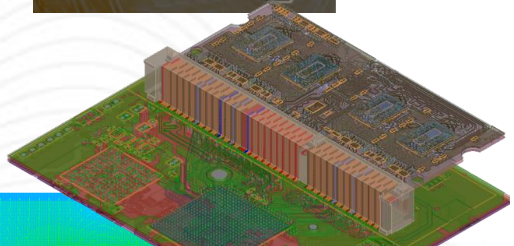

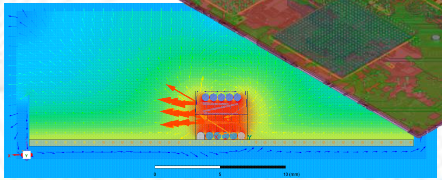

HFSS Mesh Fusion: Large Complex “EM Systems”

- Television touchscreen in EMI chamber um to meters
- Three mesh technologies in one

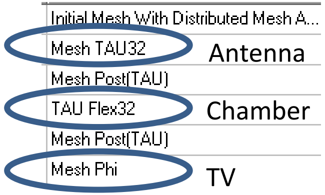

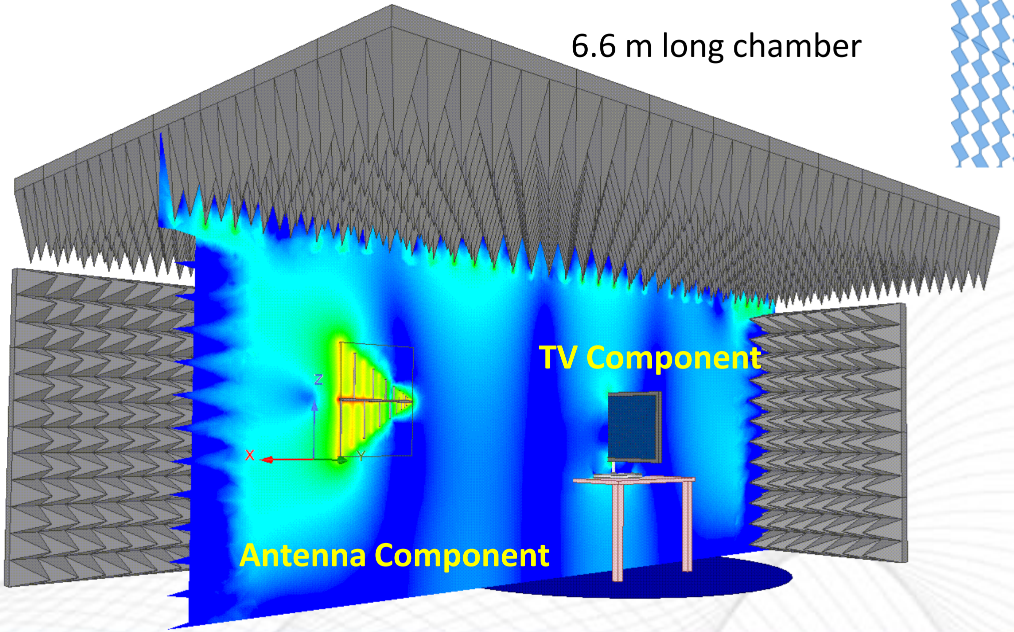

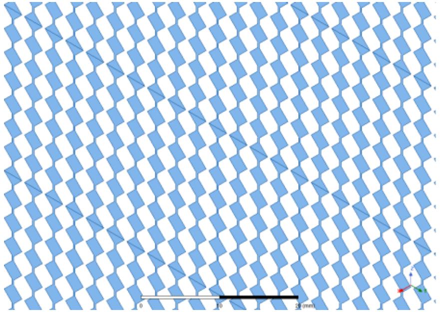

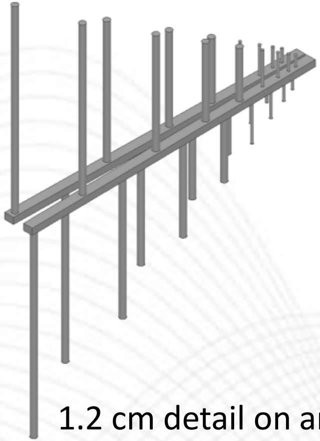

1.2 cm detail on antenna

Mesh Fusion Example: SMD Inductor on Test Board

- Unique Mesh Regions for Inductor Component and native geometry
- Fully coupled fields through mesh regions boundaries
  - Magnetic coil fields continuous from component to surrounding “native” region

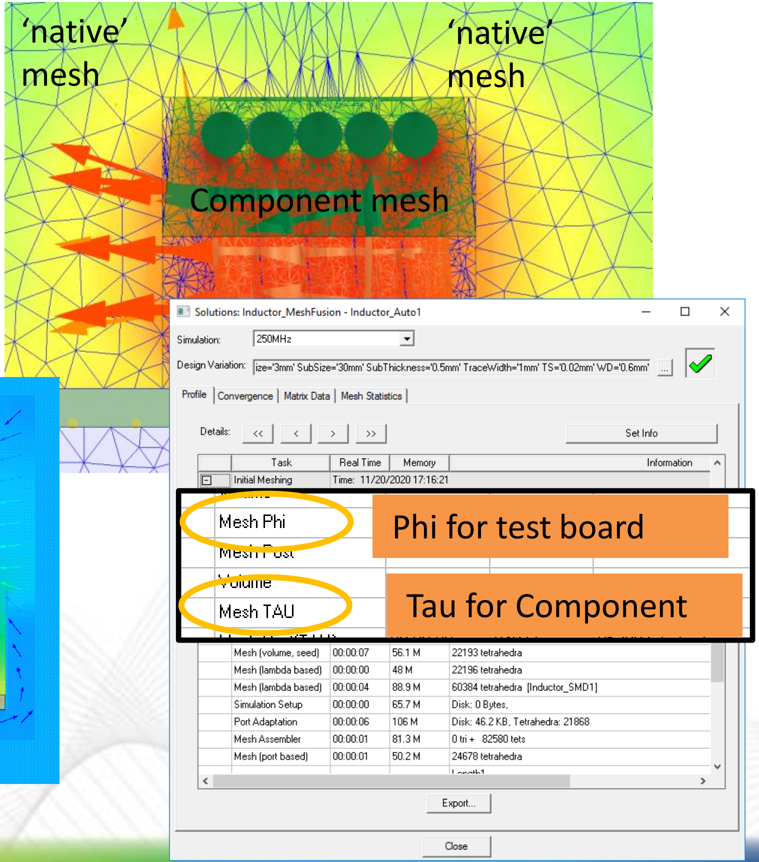

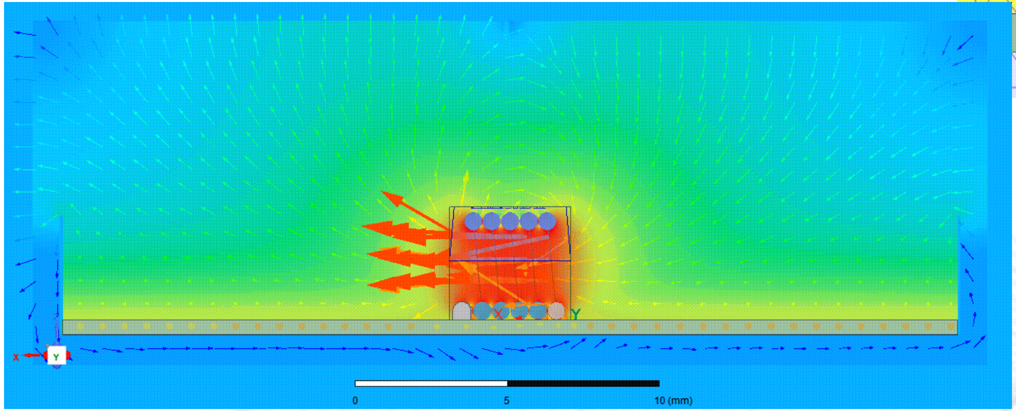

BH smoothing

Typical problems with BH curves

a) Not enough points in certain regions - need to interpolate between distant adjacent points
b) Non-smooth curves where non-physical oscillations in the 1st and 2nd derivative are observed
c) Too many points – need to reduce the number
d) Last point is too far from saturation with slope > 2 (Note if ending mu is > 100, then curve may not be able to be accurately extraploted and more data points must be input from test or datasheet)

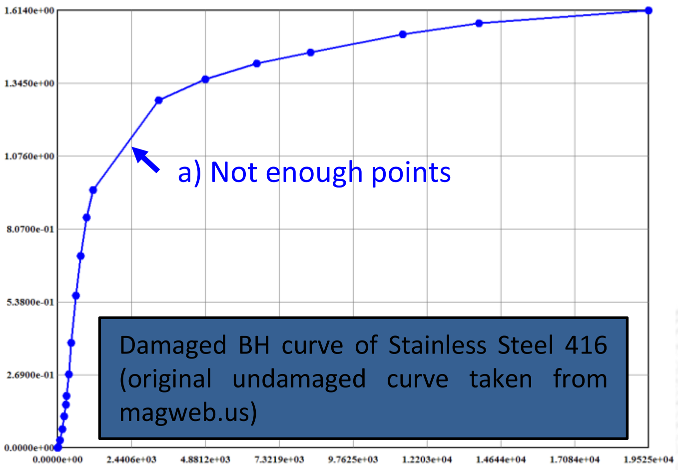

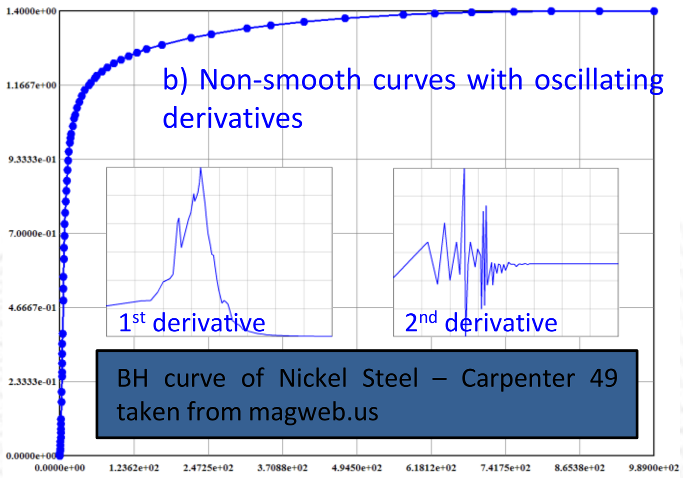

Why is BH Curve Smoothing Necessary?

It is important to have smooth BH curve and its 1st derivative

- Matrix equations are solved for components of H-field.

- For convergence Maxwell evaluates Energy Error. Energy Error is calculated from the zero divergence criteria from Maxwell's equations:

  Div B = 0 -> Div B^{approx} = $Error_n$

  $Error_n$ is the error value at the Nth element.

- Improperly defined BH curve will lead to errors. Error example:

Input data for bh curve is incorrect. For non intrinsic bh curve, the last slope of bh curve slope1 = (B_n - B_n-1)/(H_n - H_n-1) should not be larger than second last slope2 = (B_n-1 - B_n-2)/(H_n-1 - H_n-2).%1 (5:24:23 PM Dec 28, 2020)

Simulation completed with execution error on server: Local Machine. (5:24:24 PM Dec 28, 2020)

Maxwell equations for Magnetostatic analysis

$$\nabla \times \boldsymbol{H} = \boldsymbol{J}$$

$$\nabla \cdot \boldsymbol{B} = 0$$

$$\boldsymbol{B} = \mu_0 \mu_r(\boldsymbol{H}) \cdot \boldsymbol{H}$$

Maxwell 3D

Noisy or incomplete BH curves may result in convergence issues and large solution errors

BH smoothing tool

The BH smoothing tool can be used to improve BH curves **BH Curve Smoother**

Files: *.tab,*.txt and *.der

H [A/m] B [T]

Nickel-Steel-Carpenter-49.tab

H [A/m]	B [T]
0.0	0.000
0.1	0.003
0.4	0.009
0.6	0.017
0.8	0.022
1.0	0.033
1.3	0.045
1.5	0.057
1.8	0.072
1.9	0.084
2.1	0.100
2.3	0.116
2.8	0.170
3.1	0.201
3.5	0.250
3.6	0.263
3.9	0.292
4.1	0.321
4.4	0.353
4.6	0.385
5.5	0.490
5.8	0.526
6.2	0.563
6.8	0.619
7.3	0.656
7.8	0.693
8.3	0.729
9.0	0.765
9.7	0.800
11	0.834

BH curve (logarithmic view)

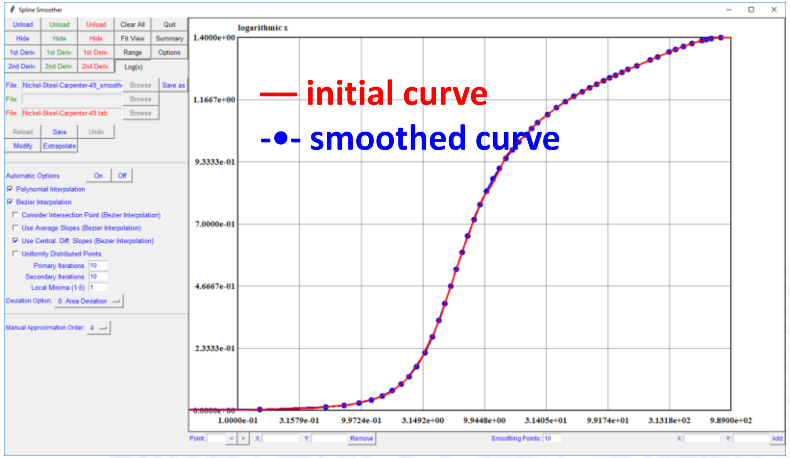

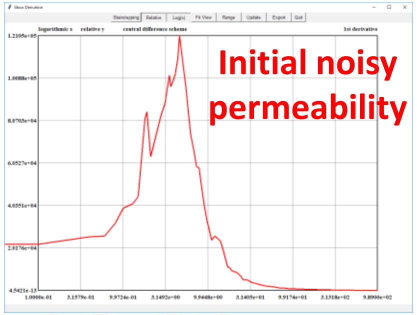

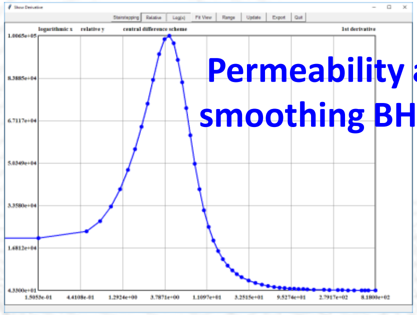

Manual and automatic smoothing techniques, Polynomial approximation, Bezier function interpolation, multiple extrapolation techniques

Meshing. Adaptive Meshing

Adaptive Meshing:

- First, Maxwell generates a solution based on a coarse initial mesh
- Then, it refines the mesh in areas of high error density and generates a new solution

Note: You can help Maxwell with specifying manually Initial Mesh setup. However, you need to be careful to not over-mesh

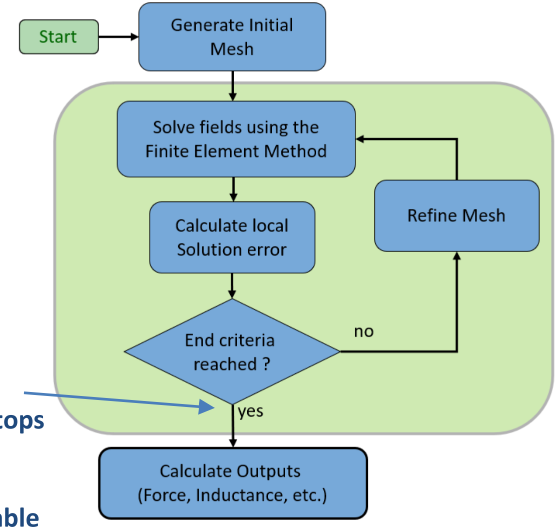

When energy % error reaches specific level adaptive refinement stops

Initial Mesh

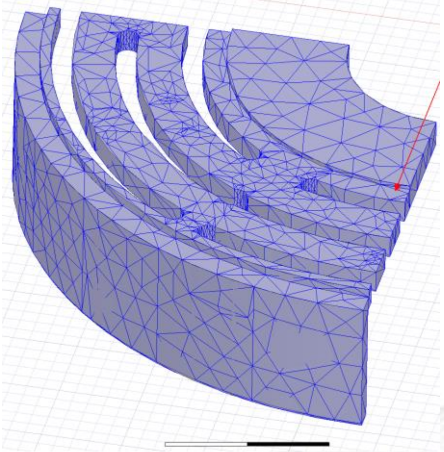

Adaptively refined Mesh (last pass)

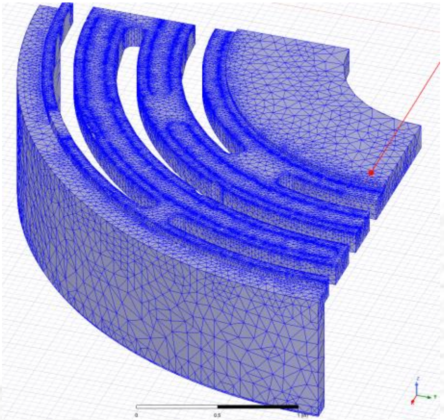

Mesh Convergence Table

Solutions: Ex_9_8_BasicOptimetrics - Puck_Attractor

Simulation: Setup1

Design Variation: move='0mm'

Profile | Convergence | Force | Torque | Matrix | Mesh Statistics

Number of Passes
- Completed 10
- Maximum 10
- Minimum 2

Energy Error/Delta Energy (%)
- Target (1, 1)
- Current (6.7657, 2.0704)

View: Table / Plot

Export...

Pass	# Tetrahedra	Total Energy (J)	Energy Error (%)	Delta Energy (%)
1	252	0.0090454	100.78	N/A
2	332	0.009181	48.79	1.4994
3	441	0.0088437	40.743	3.6748
4	579	0.0089233	38.941	0.90079
5	757	0.0088441	40.238	0.88831
6	991	0.0092404	30.405	4.4813
7	1296	0.0091963	14.802	0.47723
8	1695	0.0092247	12.23	0.30945
9	2209	0.0093259	9.309	1.0969
10	2878	0.009519	6.7657	2.0704

Meshing. Initial Mesh

Initial Mesh Settings Method: Auto, TAU, Classic

- **Auto (the default) – the solver automatically selects the mesher. In most cases, this will be TAU**
- **Classic mesher:**
  Can work better <u>on geometries with many thin, flat objects</u>. However, in some cases it <u>might not be suitable for curved surfaces</u>. Classic mesh does not have a minimum aspect ratio constraint, TAU has one.
- **TAU (Triangular Adaptive Uniform):**
  TAU has <u>many dedicated functions for true surfaces</u>. Creates more regular mesh, ideal for transient analysis where adaptive meshing is not an option

Classic

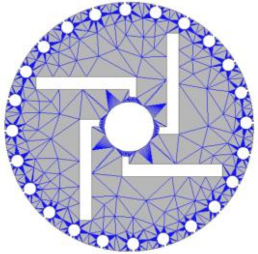

TAU

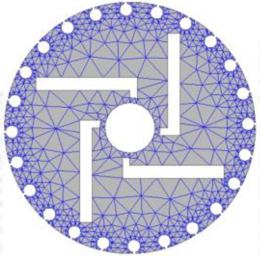

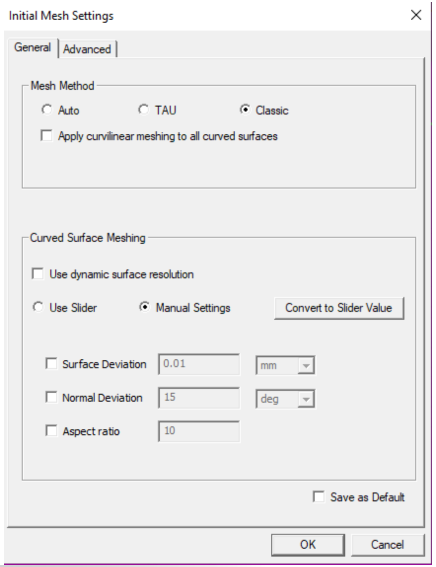

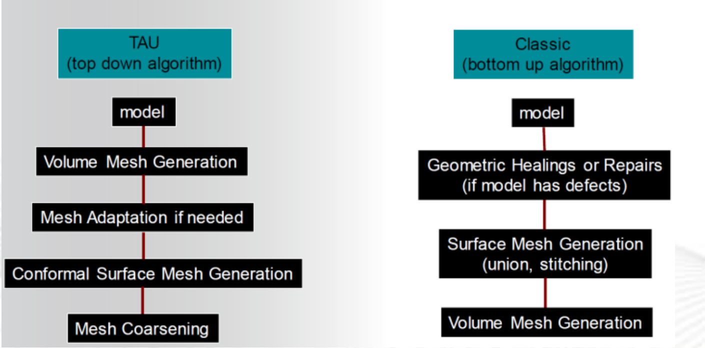

* The original model surfaces are called true surfaces

High Performance Computing

- HPC (High Performance Computing)
  - Requires additional license and it works for 3D only
  - Multiprocessing in our static solvers (MS, Eddy, ES)
  - **SDM** (Spectral Decomposition Method or Frequency sweeps) in eddy current solver.
  - Full parallelization in Transient solver with the possibility to turn **TDM** (Time Domain decomposition Method)
  - The Multi-Threading includes:
    - Initial Tau Mesh
    - Non Linear Newton-Raphson Loop
    - Matrix Assembly
    - Matrix Solving
    - Matrix Postprocessing
  - Enables to distribute parametric analysis
  - Select the menu item *Tools → Options → HPC and Analysis Options*

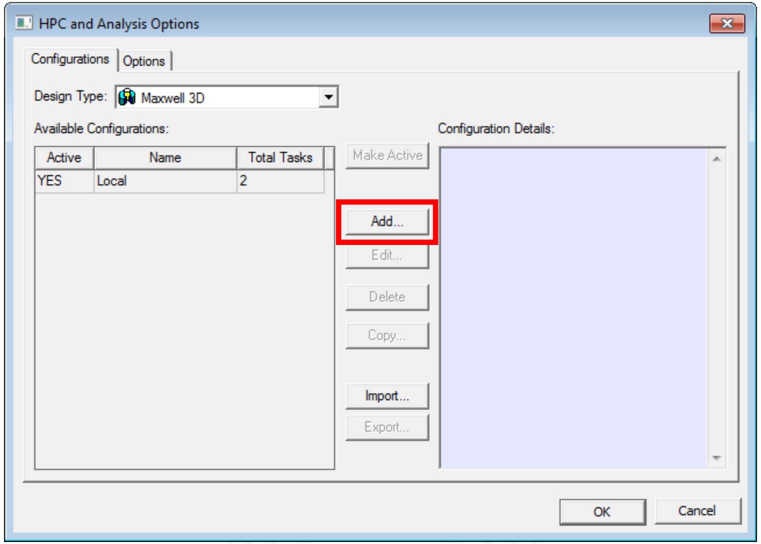

Optimetrics overview

- Optimetrics is an add-on module for ANSYS Electronics Desktop suite which provides numerous analysis tools :
  - Parametric Variation
  - Analytic Derivatives & Tuning
  - Optimization
  - Sensitivity
  - Statistical
  - Design of Experiments DOE, Response surface

- Optimetrics allows centralized control of design iterations from one common interface
- Optimetrics studies can be setup with HPC and/or TDM to parallelize runs -> can provide significant calculations speed-up

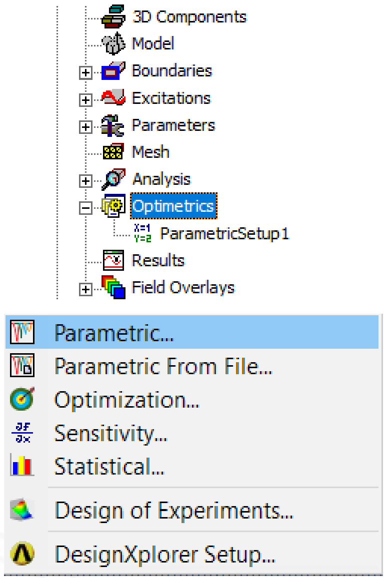

Optimetrics study process:

- Create parameterized model
- Define parameters to vary
  - Material properties, geometry, excitation, etc.
- Perform analyses
  - Parametric Sweeps
  - Optimization
  - Sensitivity Analysis
  - Statistical Analysis

Optimetrics with HPC

- **Check HPC settings for solving multiple design variations in Parallel**

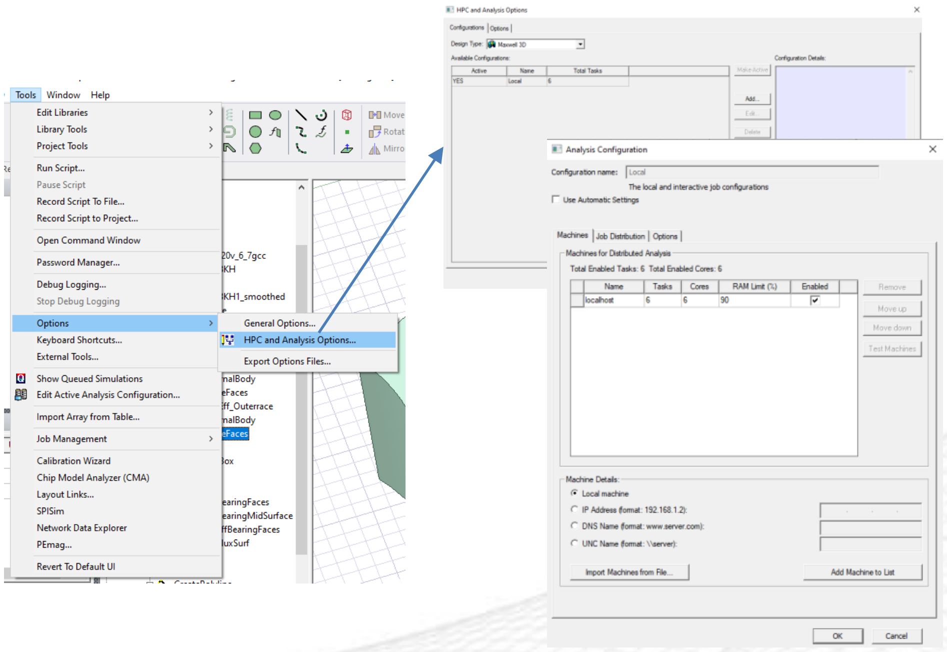

Example: solving 6 designs in parallel

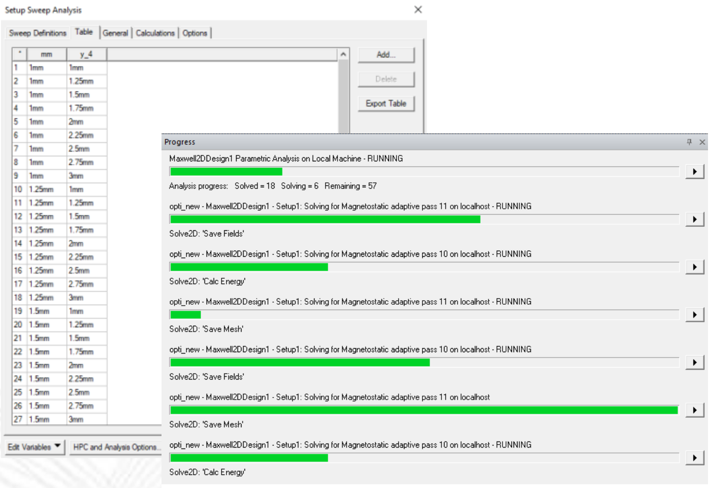

Parametric Analysis

- **General Tab**
  - **Sim. Setup:**
    - **Enables to select the required simulation setup for which parametric sweep needs to be assigned**
    - **Solver settings used in selected Simulation setup will be used to solve all design variations**

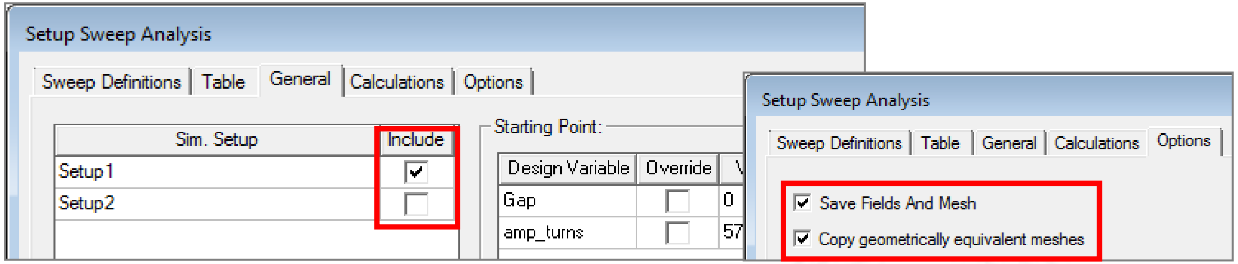

- **Options tab**
  - **Save Fields And Mesh:**
    - **Saves fields and mesh data for all the solved design variations**
    - **Design variations can be postprocessed using all postprocessing options discussed earlier**
  - **Copy geometrically equivalent meshes**
    - **Avoids remeshing if changes in input variables does not affect the geometry**

Solution Data

- **Solution Data**
  - **Solution Data contains all the information related to executed solution process**
  - **The Solution Data window can also be opened while the solution process is running to check solution convergence**
  - **Can be accessed from menu item *Maxwell 2D/3D → Results → Solution Data***

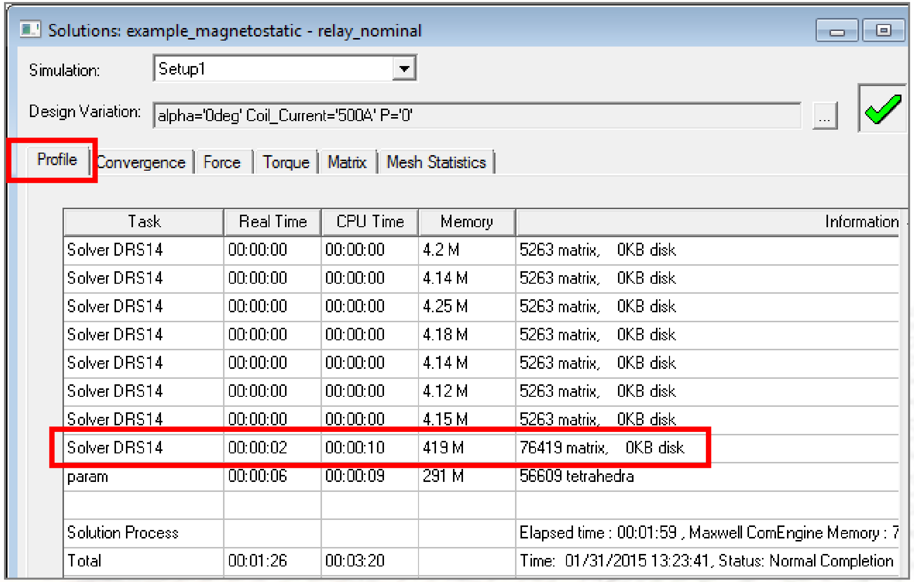

Task	Real Time	CPU Time	Memory	Information
Solver DRS14	00:00:00	00:00:00	4.2 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.14 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.25 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.18 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.14 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.12 M	5263 matrix, 0KB disk
Solver DRS14	00:00:00	00:00:00	4.15 M	5263 matrix, 0KB disk
Solver DRS14	00:00:02	00:00:10	419 M	76419 matrix, 0KB disk
param	00:00:06	00:00:09	291 M	56609 tetrahedra
Solution Process				Elapsed time : 00:01:59 , Maxwell ComEngine Memory : 7
Total	00:01:26	00:03:20		Time: 01/31/2015 13:23:41, Status: Normal Completion

- ***Profile Tab:***
  - **Contains log of tasks performed by Maxwell during solution process and time taken for each task**
  - **It reports peak physical memory used for each task**
  - **Listed tasks can be different based on type of solution being carried out**
  - **Tasks that can use HPC licenses also show the number of processors being used.**

Solution Data

- **Convergence Tab:**
  - **Reports Adaptive Convergence information**
  - **Available only with Static Solvers**
  - **Can be viewed as a Table or Plot**

Profile | Convergence | Force | Torque | Matrix | Mesh Statistics

Number of Passes
Completed 9
Maximum 10
Minimum 2

Energy Error/Delta Energy (%)
Target (1, 1)
Current (0.97, 0.24766)

View: (•) Table ( ) Plot

Export...

Pass	# Tetrahedra	Total Energy (J)	Energy Error (%)	Delta Energy (%)
1	6925	0.0073137	32.544	N/A
2	9006	0.0070874	9.9403	3.0947
3	11714	0.0069485	7.4082	1.9594
4	15233	0.0068706	5.3986	1.1221
5	19807	0.0068127	3.491	0.84167
6	25756	0.0067654	2.3529	0.69444
7	33488	0.0067348	1.8931	0.45276
8	43541	0.0067137	1.3988	0.3136
9	56609	0.006697	0.97	0.24766

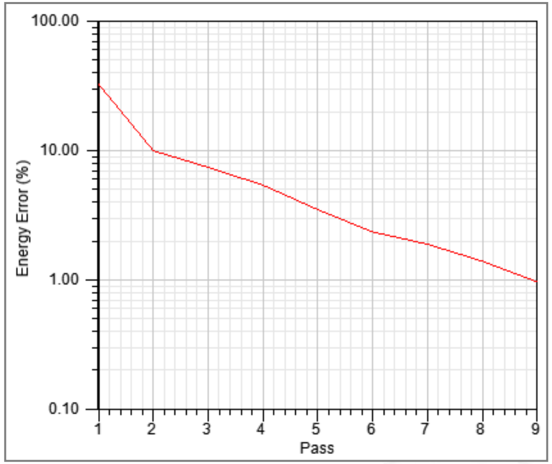

- **Force, Torque and Matrix Tab:**
  - **Reports computed parameters values**

Profile | Convergence | Force | Torque | Matrix | Mesh Statistics

Parameter: frc_arm   Force Unit: newton

Pass: 9

	F(x)	F(y)	F(z)	Mag(F)
Total	0.00070726	-0.17241	6.781	6.7832

- **Mesh Statistics Tab:**
  - **Reports mesh information and statistics**

Profile | Convergence | Force | Torque | Matrix | Mesh Statistics

Total number of elements: 56609

	Num Tet...	Min edge len...	Max edge len...	RMS edge len...	Min tet vol...	Max tet vo...	Mean tet v...
armature	4682...	0.093326...	2.02271	0.788658	1.99997e-0...	0.350564...	0.0200986...
coil	6014...	0.39737	5.56725	1.77265	0.0025895...	3.9174...	0.238649...
core	2206...	0.275647...	3.97286	1.46459	0.0009330...	2.65256...	0.0735492...
region	37894...	0.105655...	20.7723	3.84229	1.84568e-0...	400.327...	4.69778...
yoke	5813...	0.120925...	3.78255	1.1414	0.0001216...	2.34904...	0.0626543...

Batch Solver Output

See .log file in **{project name}.aedt.batchinfo** folder

- Solve starts with some job identity info, and a resource request summary

```
Machines:
HPC10 [64340 MB]: RAM: 90%, task0:8 cores, num gpus:0, gpu indexes:
HPC11 [64340 MB]: RAM: 90%, task0:8 cores, num gpus:0, gpu indexes:
HPC12 [64340 MB]: RAM: 90%, task0:8 cores, num gpus:0, gpu indexes:
```

- Next look for meshing & adaptive passes:

```
Adaptive Meshing : Time:  03/10/2021 03:34:05 (3:34:05 AM  Mar 10, 2021)
Adaptive Pass 1 : Frequency:  60GHz (3:34:05 AM  Mar 10, 2021)
…
Solver DCS8 : Real Time 00:00:04 : CPU Time 00:00:18 : Memory 1.24 G : Disk = 1.37 KB,
matrix size 254066 , matrix bandwidth  22.2
… {much later}
Adaptive Pass 10
Adaptive Passes converged
Simulation Summary:
…
Adaptive Meshing : Elapsed time: 00:02:12, total memory: 8.37 GB
max solved tets: 109152, max matrix size: 837932, max bandwidth:  51.4
```

Solver tasks add this number to show how many cores they are assigned.

Batch Solver Output

- Frequency sweep output look like this:

```
Frequency - 65.625GHz on AlexDesktop.simutechgroup.ca
Simulation Setup : Real Time 00:00:01 : CPU Time 00:00:01 : Memory 142 M : Disk = 0 Bytes, 67394 tetrahedra
Matrix Assembly : Real Time 00:00:05 : CPU Time 00:00:08 : Memory 490 M : Disk = 0 Bytes, 67394 tetrahedra , 1: 101 triangles , 2: 112 triangles
Solver DCS2 : Real Time 00:00:24 : CPU Time 00:00:43 : Memory 949 M : Disk = 0 Bytes, matrix size 428768 , matrix bandwidth  22.2
```

- Maxwell Time Domain Tasks look like this:

```
Solve TD8 : Real Time 00:01:44 : CPU Time 00:09:56 : Memory 8.91 G : 817550 tetrahedra
Solver Progress : Completed time point 0.0148333s
Solver DRS8, 4 iterations : Real Time 00:00:15 : CPU Time 00:01:48 : Memory 1.36 G : 380303 matrix,  0KB disk
Solve TD8 : Real Time 00:01:46 : CPU Time 00:09:59 : Memory 8.9 G : 817550 tetrahedra
Setup1 : [PROFILE] Solver Progress : Completed time point 0.01425s
Solver DRS8, 4 iterations : Real Time 00:00:15 : CPU Time 00:01:51 : Memory 1.41 G : 380303 matrix,  0KB disk
Solve TD8 : Real Time 00:01:44 : CPU Time 00:09:55 : Memory 8.89 G : 817550 tetrahedra
```

- Shows cores, memory usage, elements, solve time, and more, but making predictions based on text output is rather difficult.

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+

Mechanical Test Data:

- 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.

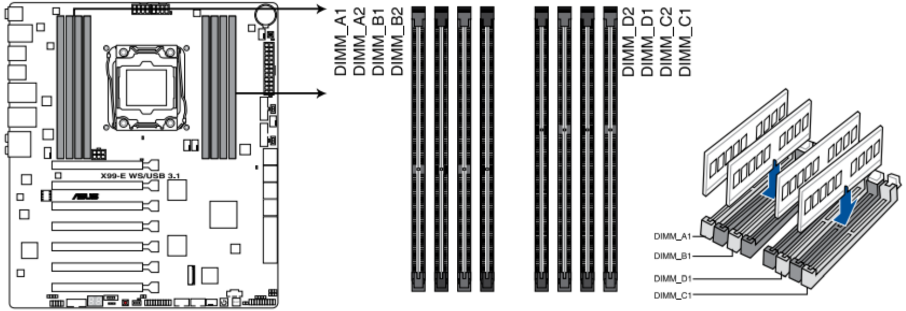

X99-E WS/USB 3.1 288-pin DDR4 DIMM socket

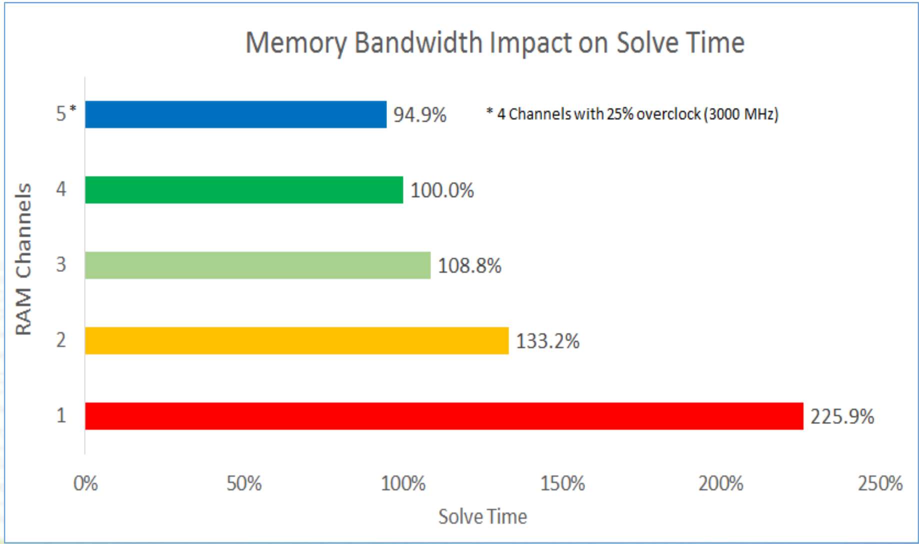

Which CPU Then?

- The direct solver is really important; 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	Example Recommended Processors
4-8	i9-10900X Xeon W Series Ryzen 5800X
12-16	i9-10920X / i9-10980XE Xeon W Series Single Xeon Gold+ Threadripper 3960X
24+	Dual Xeon Gold+ (like 6242R) Dual AMD EPYC

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

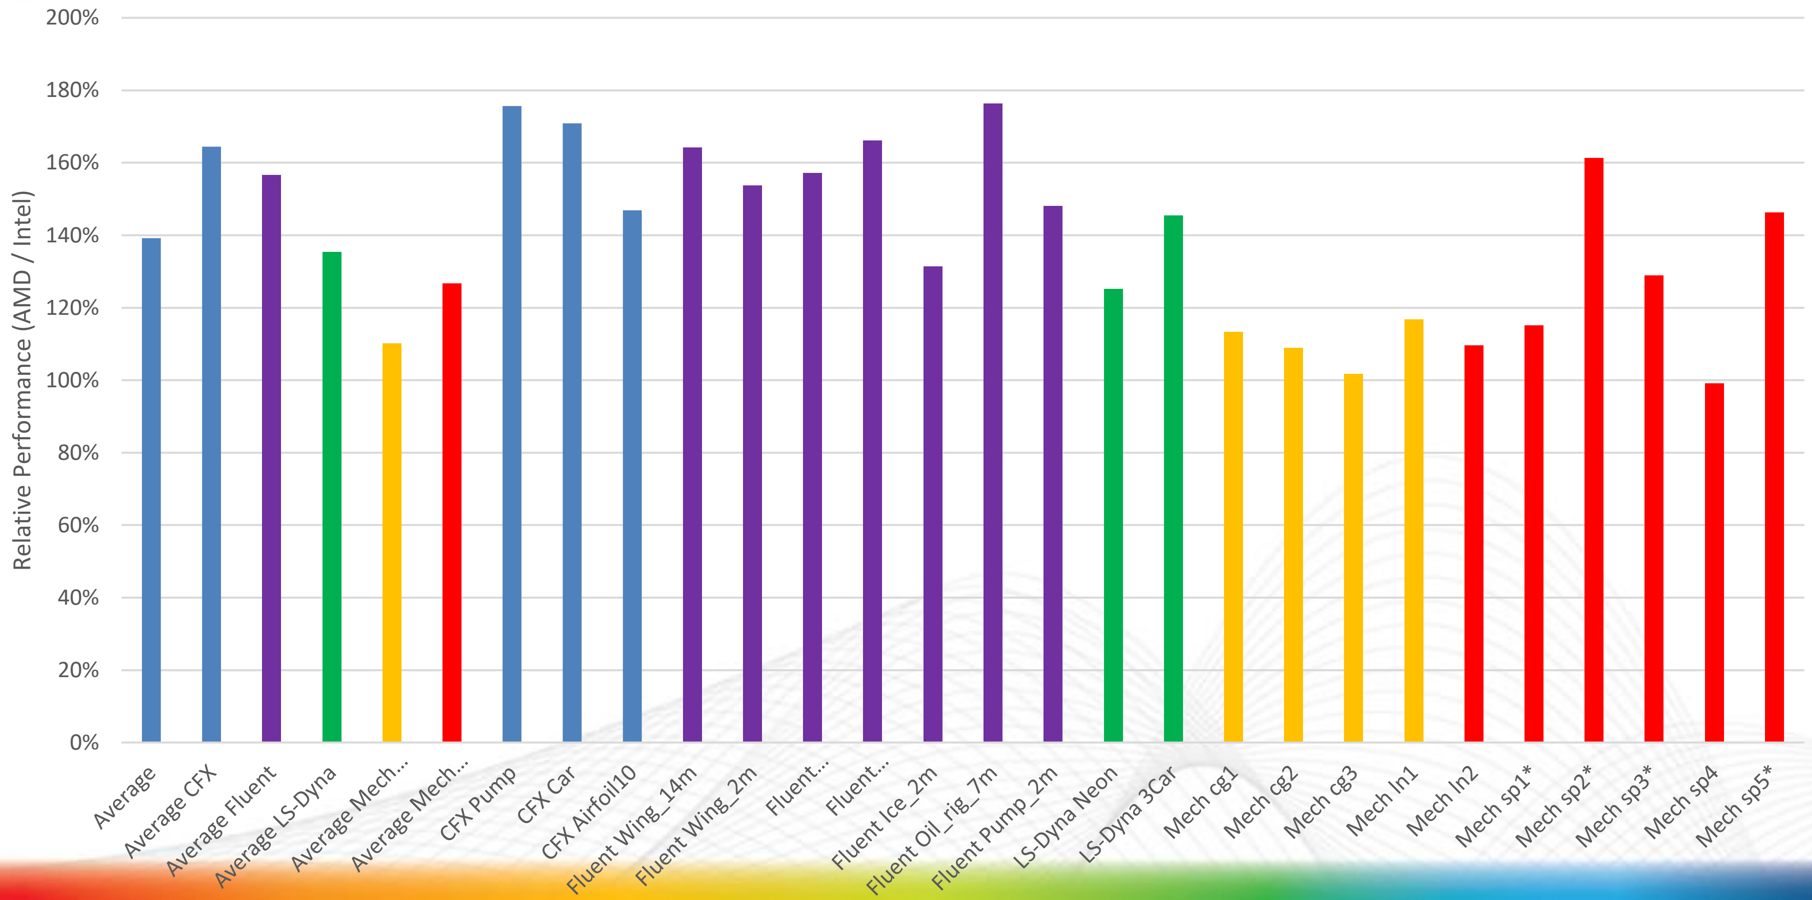

Dual EPYC 7502 vs Dual Xeon 6242R
36 cores vs 36 cores

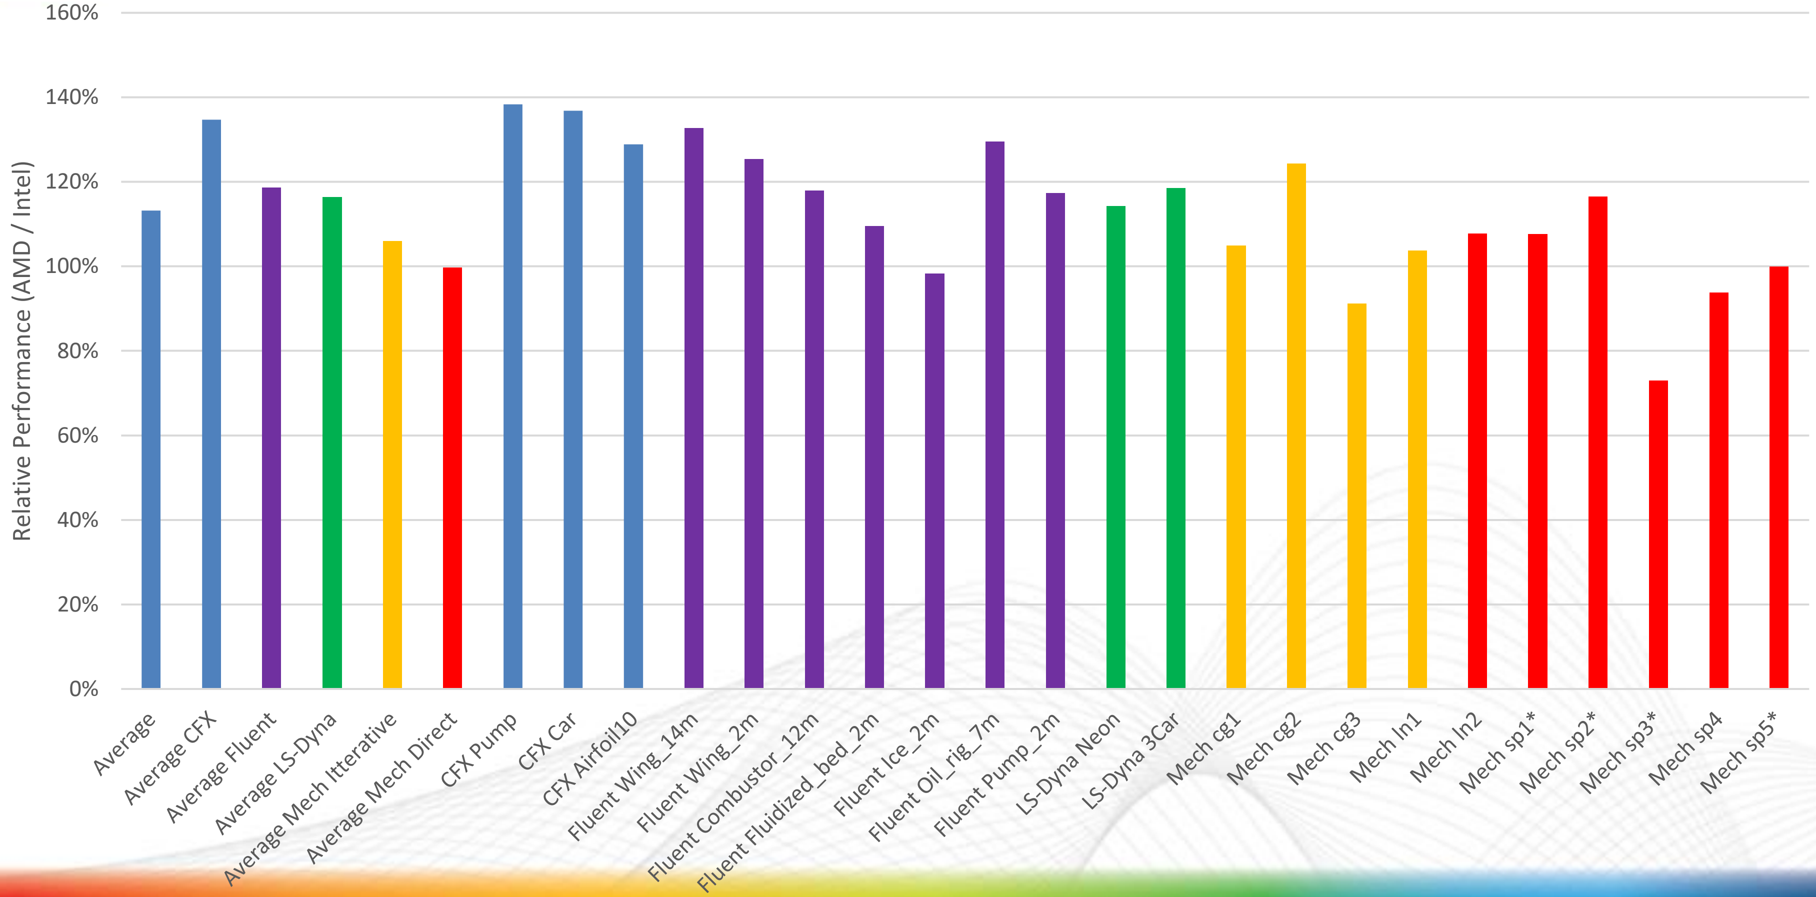

Dual EPYC 7502 vs Dual Xeon 6242R
36 cores vs 36 cores

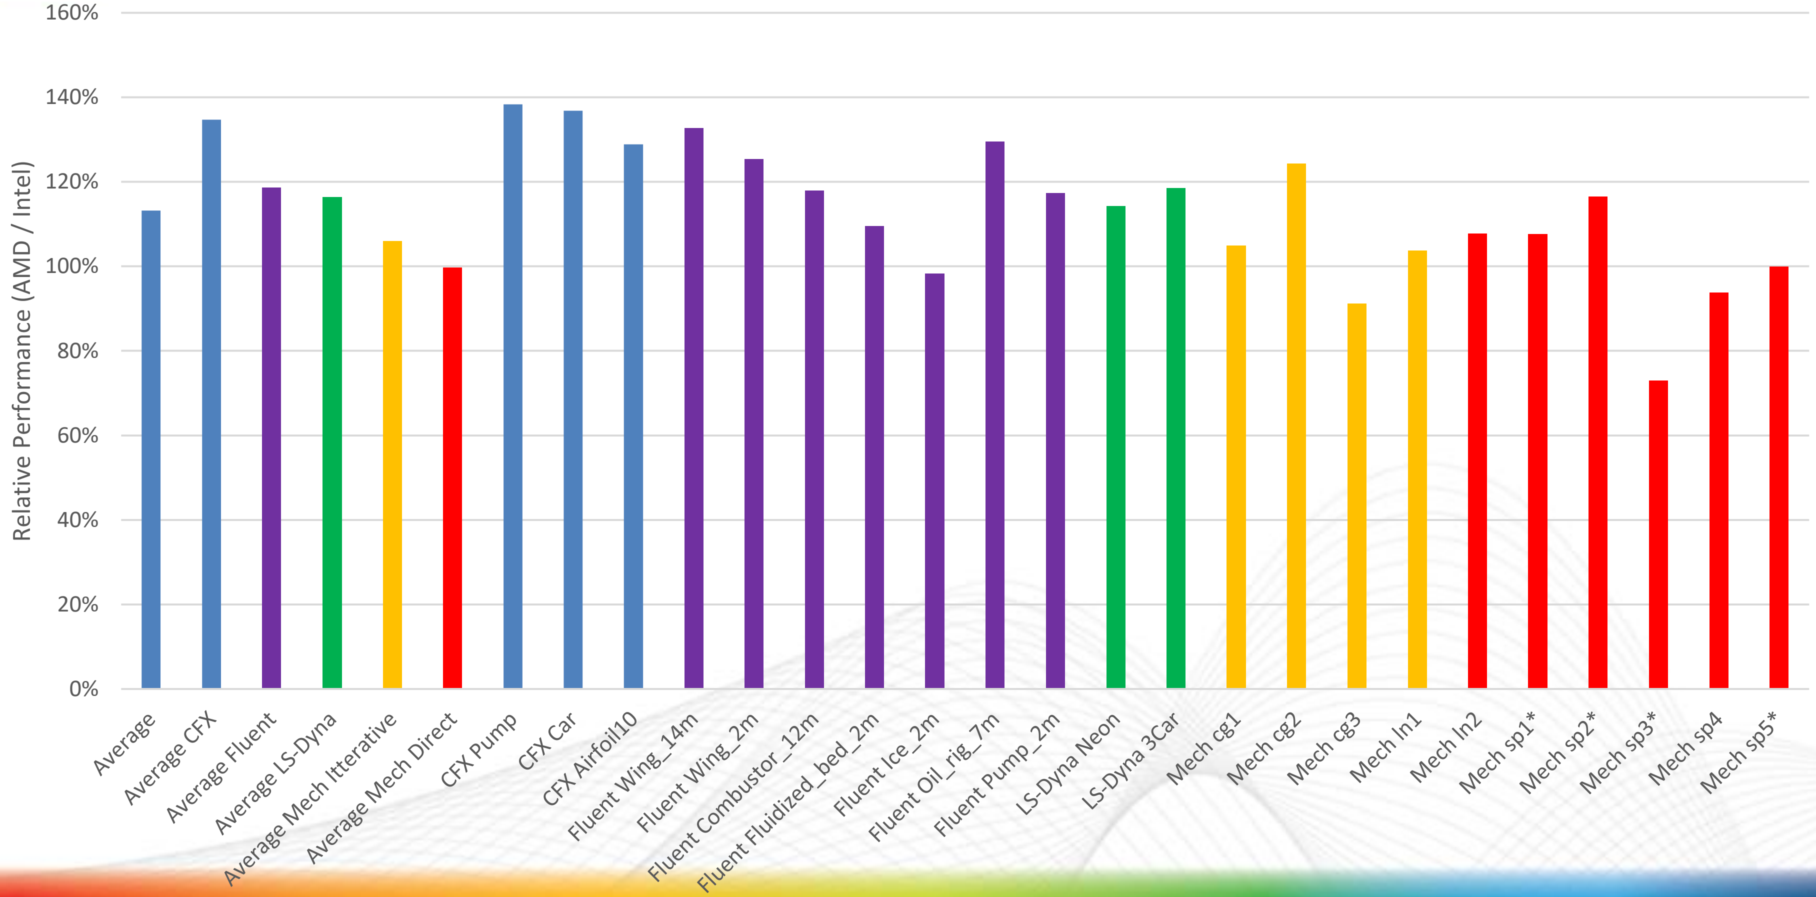

Threadripper 3960X vs i9-9960X
12 cores vs 12 cores Mechanical Tests

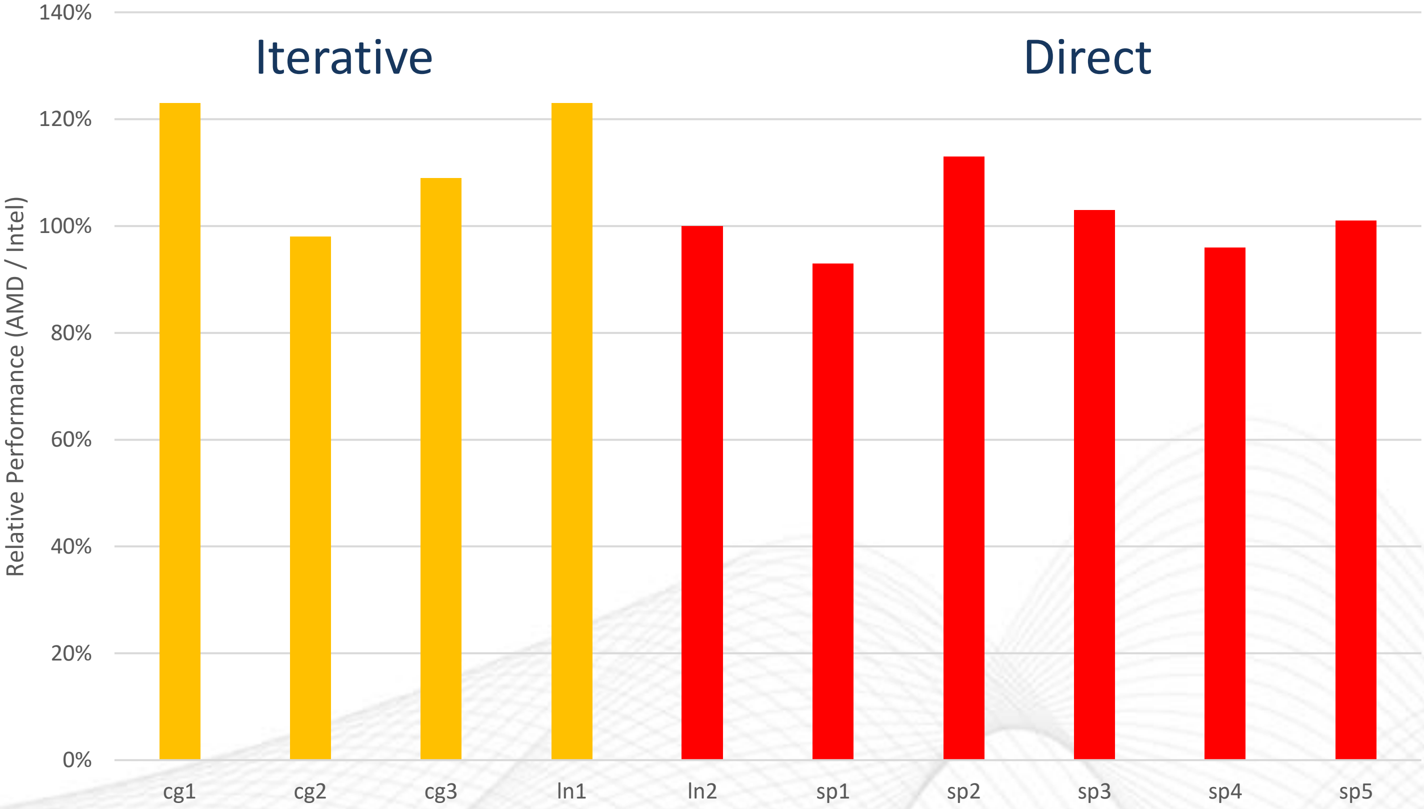

Threadripper 3960X vs i9-9960X
24 cores vs 16 cores Mechanical Tests

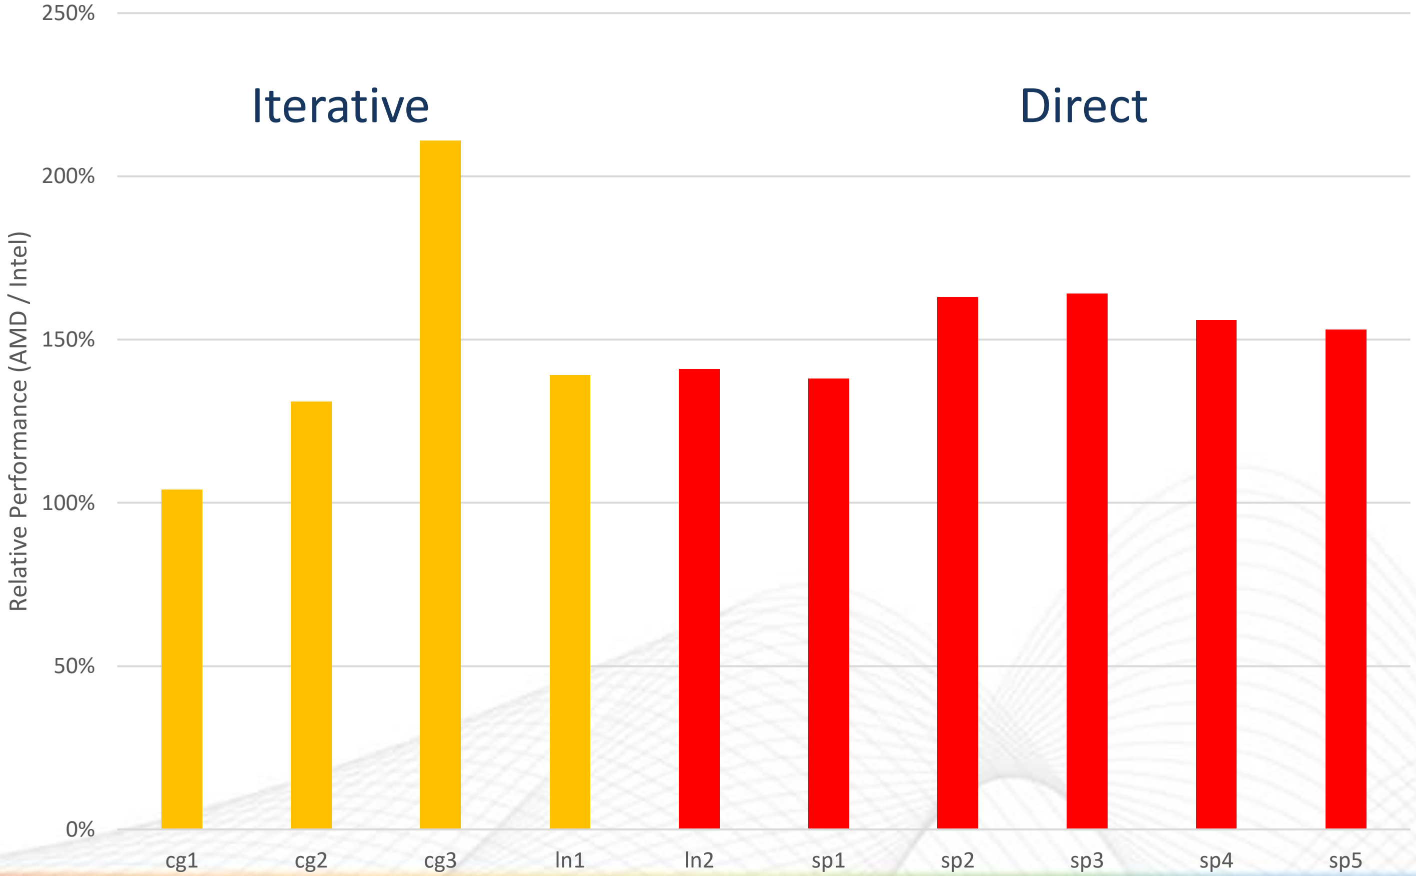

Regarding RAM Sizing

- Giving a memory recommendation is difficult because it is very usage dependent, driven by fidelity requires, **platform limits**, and budget.
- Extra RAM can be very helpful as it increases the parallelization available using the best techniques.
- Sizing for 100% of use cases is impractical, aim to get an amount you will regularly benefit from.
- Largely comes down to economics and appetite, RAM is fairly cheap.
- Using a cluster to scale RAM allows much more dynamic range, and can also scale cores.
    - Need more RAM: request more cluster nodes.
- Common recommendations: 128-384 GB per compute node, sometimes more.

RAM Management

- If you don’t have enough RAM, how can you adapt:
  - Switching to iterative
  - Reduce mesh (geometry simplification, periodicity, symmetry)
  - Use different mesh type
  - Assign more cores per task (less parallel tasks)
  - Subdivide domain more (Maxwell TDM)
  - Use less tasks per computer, as little as 1 (undersubscribe resources)

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 EDT works very well, but has mixed requirements.
- Each node adds more cache, memory bandwidth, memory quantity, cores, etc.
- High speed interconnect (40gbps+) is required for nodes working on a highly dependent solutions like Domain Decomposition Method (DDM).
- RDMA communication is highly recommended for high dependency
  - Traditionally infiniband was recommended, but ethernet has this too.
  - ~2 microsecond vs 30 microsecond latency, higher bandwidth, less overhead
- Low speed interconnect (1gbps) is sufficient for independent tasks (Spectral Decomposition, DSO).

Quick Summary for Icepak

- Under the hood the Icepak solver is basically fluent, so CFD rules apply.
  - See CFD presentation slides for Fluent advice. Available with this slides after event.
- It's highly scalable, and runs best with maximum memory bandwidth and cache (AMD platform currently advantageous).
- Keep elements/core above 5-10k for maximum performance.

Questions?

Thanks for listening!

Alex Pickard (HPC)	apickard@simutechgroup.ca
Zoe Zhu (HF)	zzhu@simutechgroup.ca
Natalya Vorobtsova (LF)	nvorobtsova@simutechgroup.com

www.SimuTechGroup.com/Offices