

Forum TERATEC

20 - 21 OCTOBRE 2026

PARIS · PALAIS DES CONGRÈS

New mixed solver leveraging the power of next-generation GPUs for FEA simulation

Bruno Gaudin – Ansys part of Synopsys
bruno.gaudin@synopsys.com

Agenda

- FEA simulation and context
- Hardware CPU / GPU market trend
- 3 categories of GPU
- Mixed solver : an “in-between” solver
- Performance and memory footprint
- Conclusion

FEA simulation and context

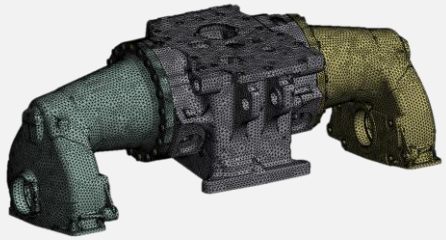
- Main goal: **to predict by calculation the behavior of a structure under mechanical, thermal loads**

- **Strategy:**

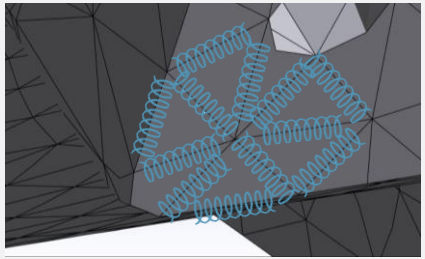
1. *To divide the model in millions of small elements (meshing). Individually each element behaves likes some springs.*
2. *To characterize each spring with correct equivalent stiffness*
3. *To assembly of one BIG stiffness matrix K*
4. ***To invert K matrix***
5. *Get results*

- Multiple TFLOPs Required **With Double Precision**
- Main task of the workflow
- Demanding Extreme Compute Performance and Scalability

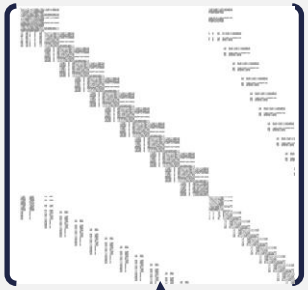
1



2



3

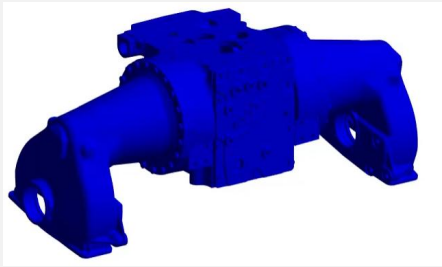


$F=K.U$

4

Matrix inversion

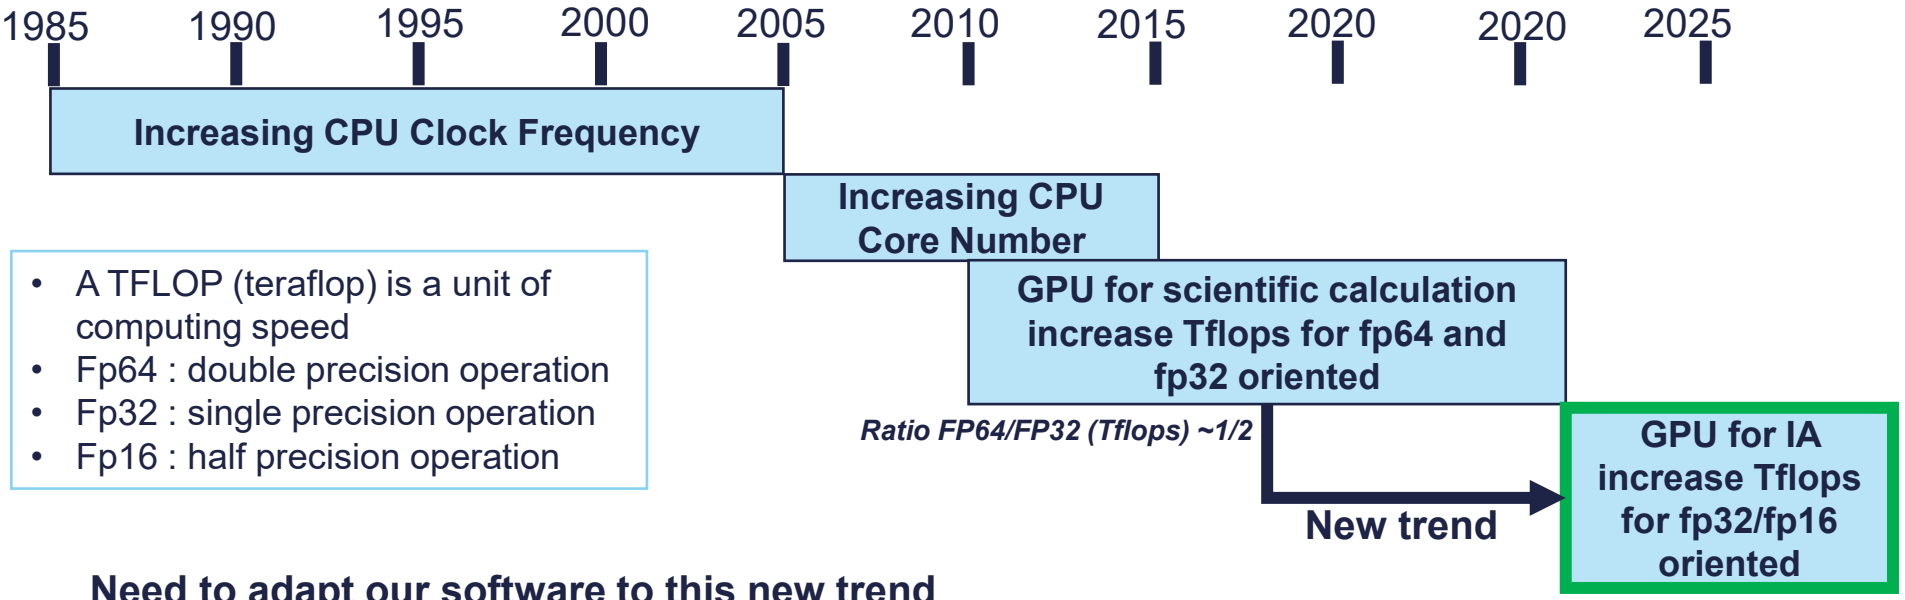
5



$U=K^{-1}*F$

Hardware CPU / GPU market trend

- Hardware evolution as performance gain for FEA simulation



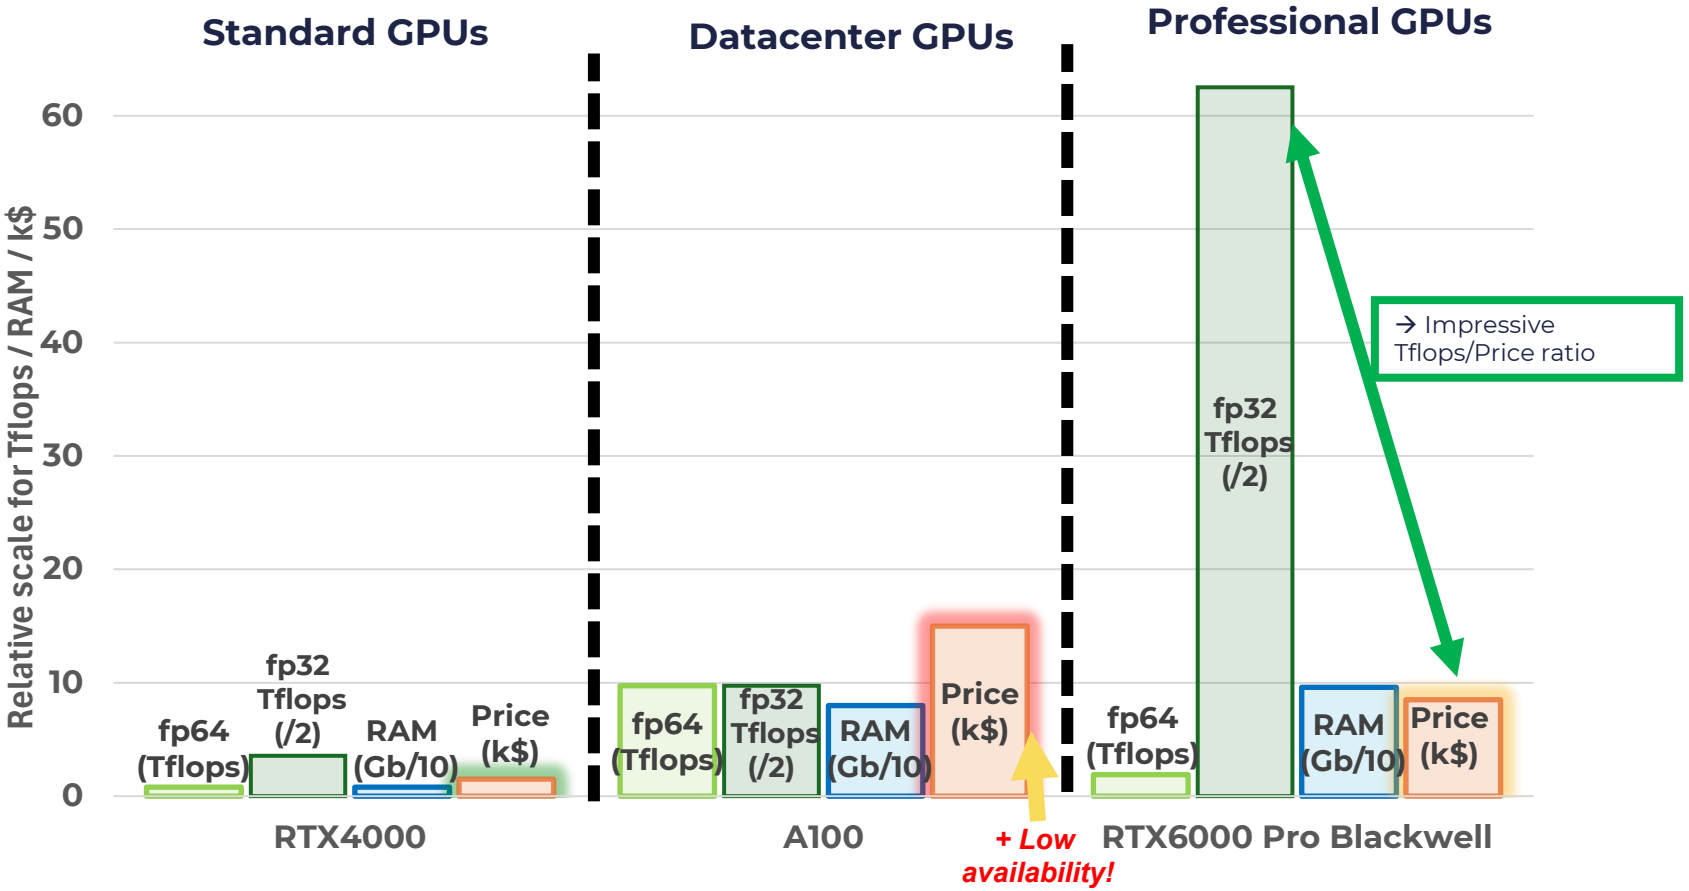
- A TFLOP (teraflop) is a unit of computing speed
- Fp64 : double precision operation
- Fp32 : single precision operation
- Fp16 : half precision operation

Need to adapt our software to this new trend
Goal: fp32 performance with fp64 final accuracy

MIXED SOLVER

3 categories of GPU (3 examples)

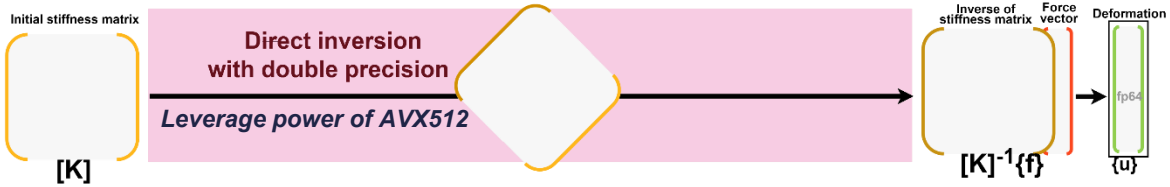
Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS



Mixed solver : an “in-between” solver

Double precision = fp64
Single precision = fp32

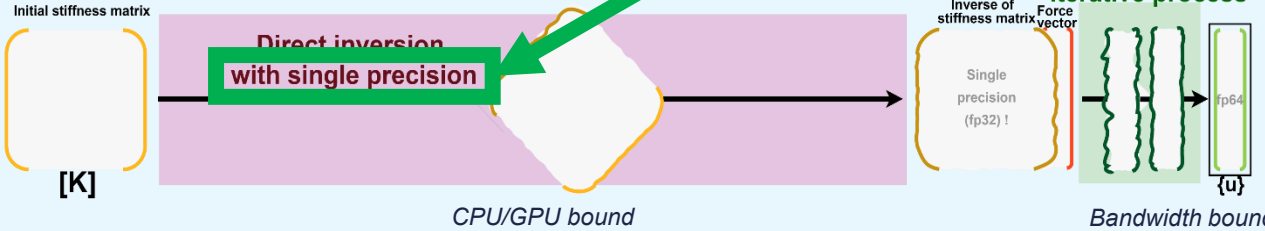
Sparse:



Mostly used and widely used

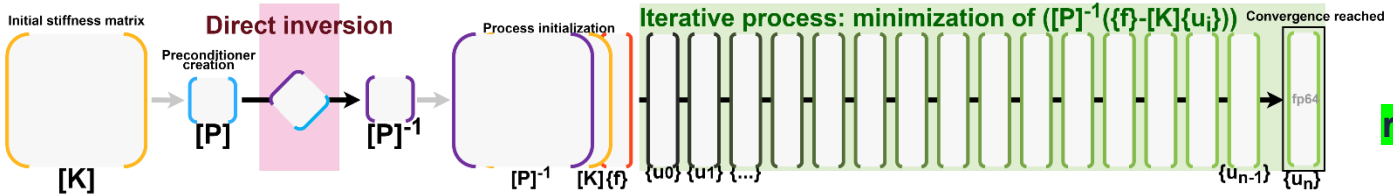
Mixed:

Leverage the power of fp32 GPU !



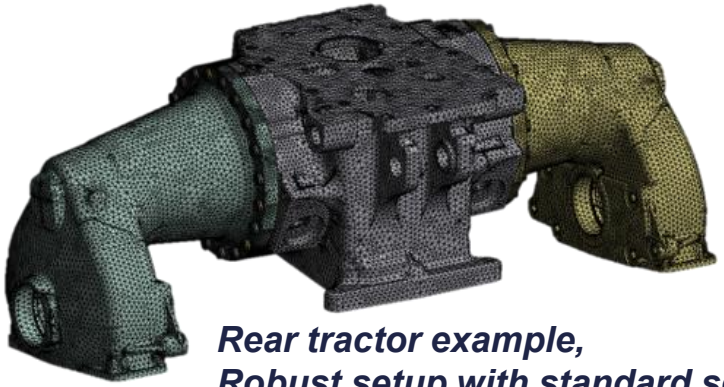
MIXED approaches → Very similar to Sparse

PCG :



Very large model relative to available memory

Mixed solver study on a robust baseline



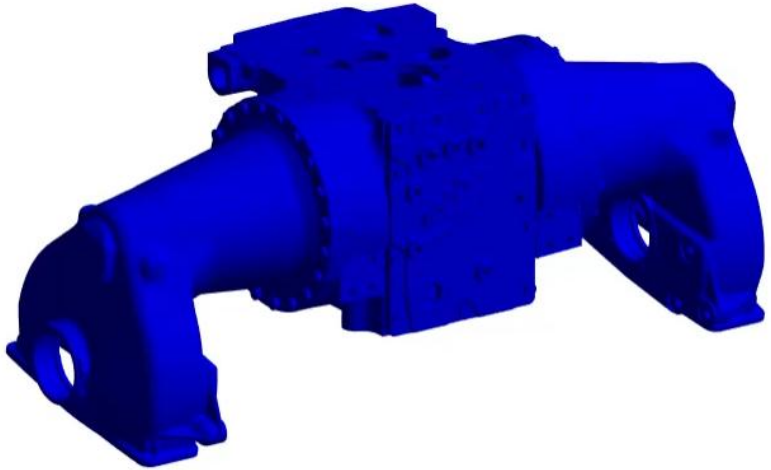
Rear tractor example,
Robust setup with standard solver

- **10 Millions degree of freedom** (matrix size) over **186 iterations** to simulate full loading convergence.

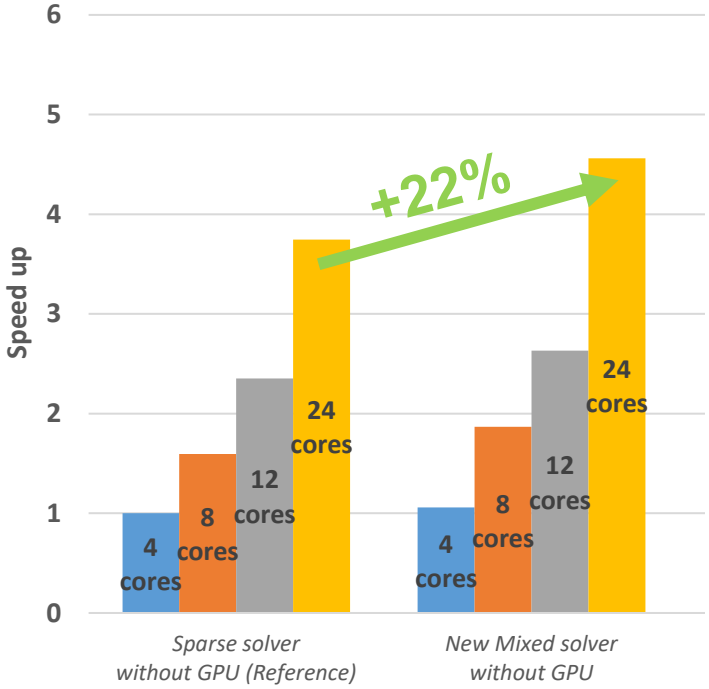
3 types of non-linearities:

- Material
- Contact
- Large deformation

→ Requires several matrix inversions (186)



Performance and memory footprint

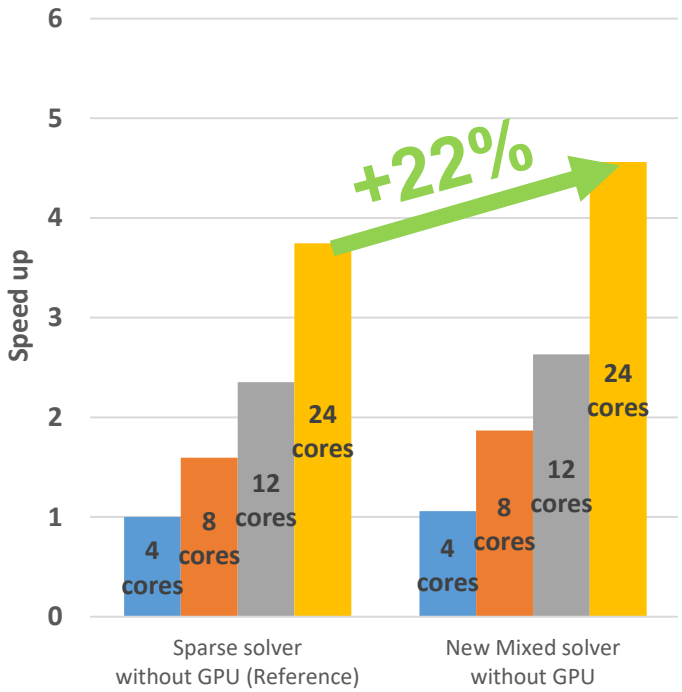


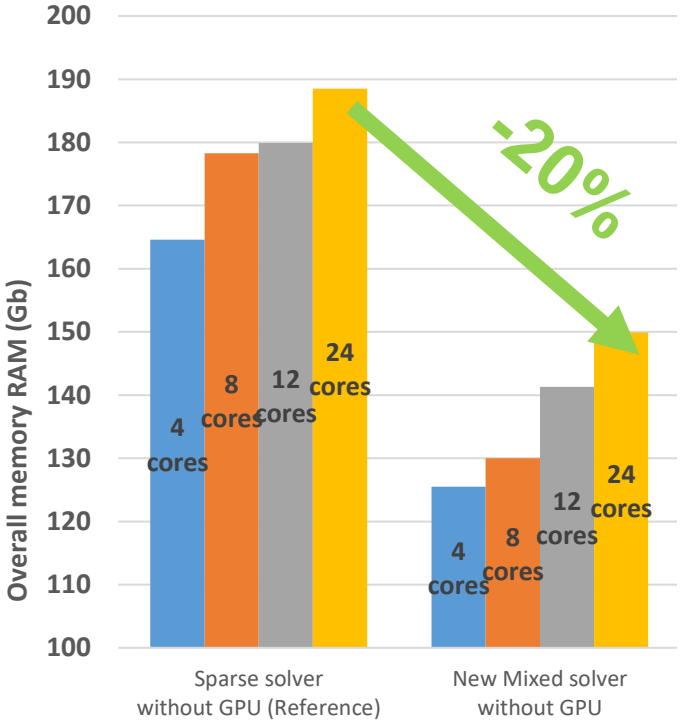
Key message 1 :

- **New mixed solver offers similar or better performance compared to Sparse solver**

Ref : server dual AMD 9575F – Dual RTX 6000 Pro Blackwell

Performance and memory footprint



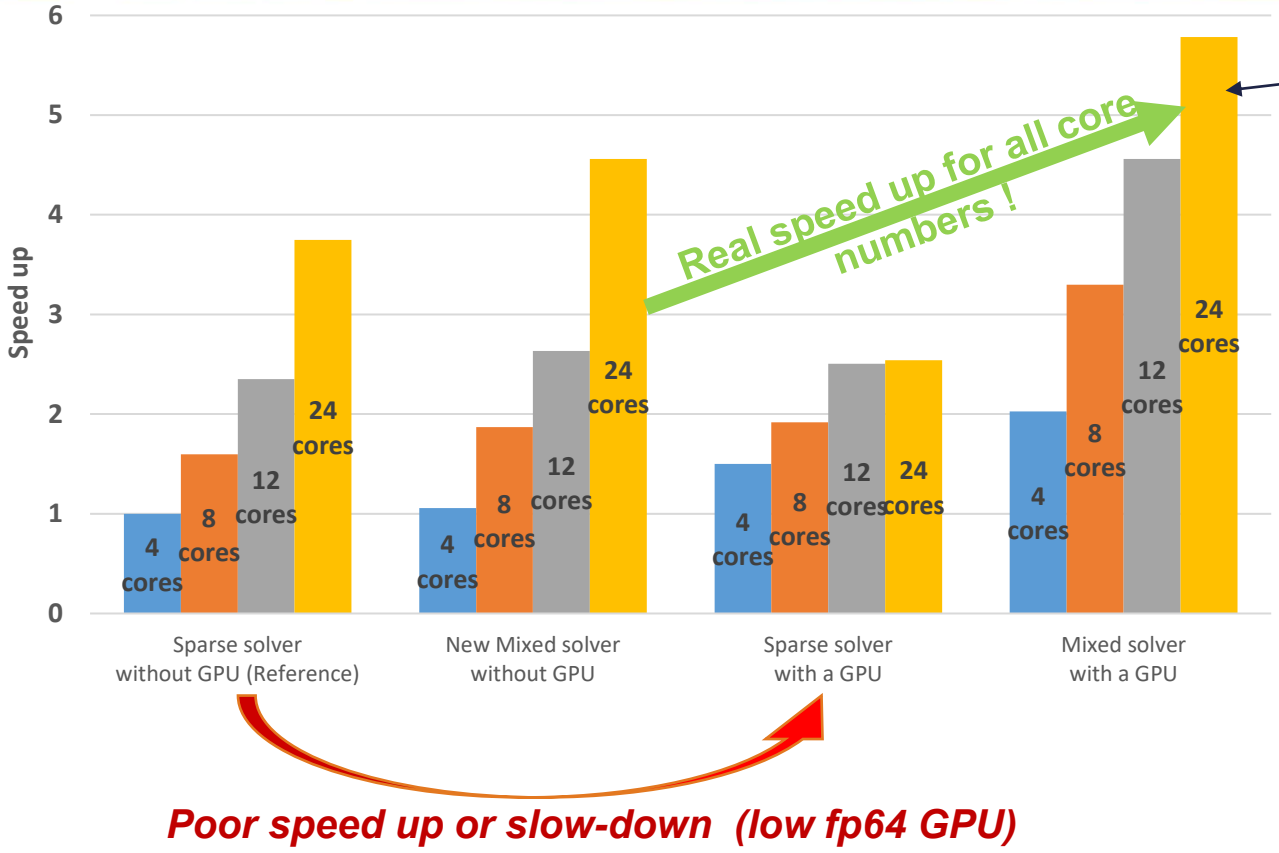


Key message 2 :

- **New mixed solver offers a better memory footprint**

Ref : server dual AMD 9575F – Dual RTX 6000 Pro Blackwell

Performance and memory footprint with GPU



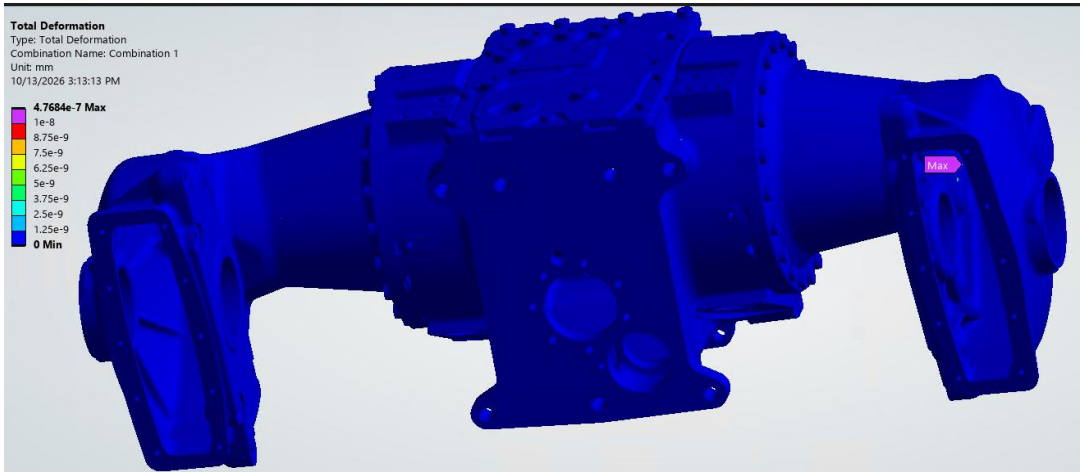
Key message 3 :

- **Mixed solver leverage the power of fp32 GPU acceleration leverage**

Ref : server dual AMD 9575F – Dual RTX 6000 Pro Blackwell

Relative displacement field difference

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS



Max : 4.77 10^-7 mm

Key message 4 :

- **Same results (same linear path) for all cases (16 runs)**

Ref : server dual AMD 9575F – Dual RTX 6000 Pro Blackwell / Ansys APDL v261 / IntelMPI2021.16 / Windows Server22

```
 --- S O L U T I O N   S U M M A R Y --------

 Analysis Type: Static (Steady-State)
 Nonlinear: Yes

                        Number
   Load Steps                4
   Iterations              186
   Bisections                0
```

Ex 1: Case 24 cores/ Mixed with GPU

```
 --- S O L U T I O N   S U M M A R Y --

 Analysis Type: Static (Steady-State)
 Nonlinear: Yes

                        Number
   Load Steps                4
   Iterations              186
   Bisections                0
```

Ex 2: Case 24 cores/ Sparse with GI

Time		
<COLDATA>		
0.5000000E-01	0.1284439E-02	0.1284439E-02
0.1000000	0.2623818E-02	0.2623818E-02
0.1750000	0.4536059E-02	0.4536059E-02
0.2875000	0.6068721E-02	0.6068721E-02
0.4000000	0.5173719E-02	0.5173719E-02
0.5687500	0.4618021E-02	0.4618021E-02
0.8187500	0.3894079E-02	0.3894079E-02
1.000000	0.3296387E-02	0.3296387E-02
2.000000	0.3130404E-02	0.3130404E-02
2.100000	0.2634338	0.2634338
2.200000	0.5285455	0.5285455
2.300000	0.8092804	0.8092804
2.400000	1.119994	1.119994
2.500000	1.544730	1.544730
2.600000	2.044421	2.044421
2.700000	3.020106	3.020106
2.800000	3.440316	3.440316
2.900000	5.384976	5.384976
3.000000	12.09299	12.09299
3.100000	11.81504	11.81504
3.200000	11.53529	11.53529
3.350000	11.11168	11.11168
3.550000	10.53507	10.53507
3.750000	9.940028	9.940028
3.875000	9.554969	9.554969
4.000000	9.134985	9.134985
</COLDATA>		
</SOLUTION>		

Ex:3 UZ displacement Case 24 cores/ Mixed with GPU over time step

Ex:4 UZ displacement Case 24 cores/ Sparse with GPU over time step

Mixed solver robust across CPU and libraries

	CPU : Intel	CPU : AMD
MKL library	186 iterations (default)	186 iterations
AOCL library	186 iterations	186 iterations (default)

```
6139 Compiler: Intel(R) Fortran Compiler Classic Version 2021.9  (Build: 20230302)
6140           Intel(R) C/C++ Compiler Classic Version 2021.9  (Build: 20230302)
6141           Intel(R) oneAPI Math Kernel Library Version 2024.2.3-Product Build 20250403
6142
6143 Number of machines requested            :    1
6144 Total number of cores available         :   48
6145 Number of physical cores available      :   48
6146 Number of processes requested           :   12
6147 Number of threads per process requested :    1
6148 Total number of cores requested         :   12 (Distributed Memory Parallel)
6149 MPI Type: INTELMPI
6150 MPI Version: Intel(R) MPI Library 2021.16 for Windows* OS
6151
6152
6153 Number of GPUs available                :    2
6154 Number of GPUs requested               :    1
6155 GPU Acceleration: NVIDIA Library Requested and Enabled
6156 GPU Device with ID =  0 is: NVIDIA A100 80GB PCIe
6157 GPU Driver Version: 13.0
6158 CUDA Version: 13.0
```

```
6139 Compiler: Intel(R) Fortran Compiler Classic Version 2021.9  (Build: 20230302)
6140           Intel(R) C/C++ Compiler Classic Version 2021.9  (Build: 20230302)
6141           AOCL-BLAS 5.1.1 Build 20251008
6142           Intel(R) oneAPI Math Kernel Library Version 2024.2.3-Product Build 20250403
6143
6144 Number of machines requested            :    1
6145 Total number of cores available         :   48
6146 Number of physical cores available      :   48
6147 Number of processes requested           :   12
6148 Number of threads per process requested :    1
6149 Total number of cores requested         :   12 (Distributed Memory Parallel)
6150 MPI Type: INTELMPI
6151 MPI Version: Intel(R) MPI Library 2021.16 for Windows* OS
6152
6153
6154 Number of GPUs available                :    2
6155 Number of GPUs requested               :    1
6156 GPU Acceleration: NVIDIA Library Requested and Enabled
6157 GPU Device with ID =  0 is: NVIDIA A100 80GB PCIe
6158 GPU Driver Version: 13.0
6159 CUDA Version: 13.0
```

Key message 5 :

- **Same results for all 4 cases here.**

Conclusion

- New mixed solver :
  - Batch solution : similar timings to complete simulation runs
  - Less memory usage
  - **Same non-linear behavior, equal convergence paths, equal results**

  - Ability to leverage fp32 power of new GPU

- In the future, mixed solver will support more analysis type (linear dynamics …)