A Subzone-Based In-Situ Technique for I/O Efficient Analysis and Visualization of Overset Grid Results

Scott Imlay, CTO
Craig Mackey, Senior Research Engineer
Scott Fowler, Product Manager

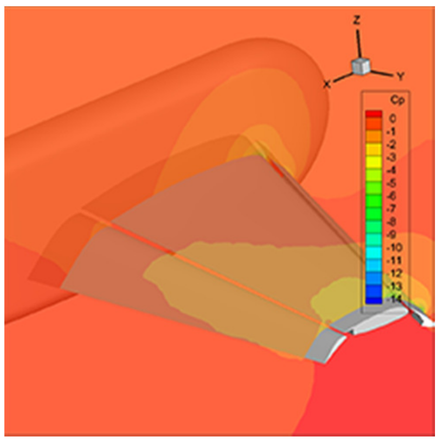

tecplot®
Master the View

Agenda:

1. Introduction & Motivation
2. SZL Technology Explained (Enabling Technology)
3. New InSitu Technique Described
4. Comparison with Traditional InSitu
5. Results

SZL Technology:

Motivation

- Wide range in length scales
- Resolution of grid
  (# of grid points)
  constrained by computer
  performance
  (growing with Moore’s law)

SZL Technology:

CFD 2030

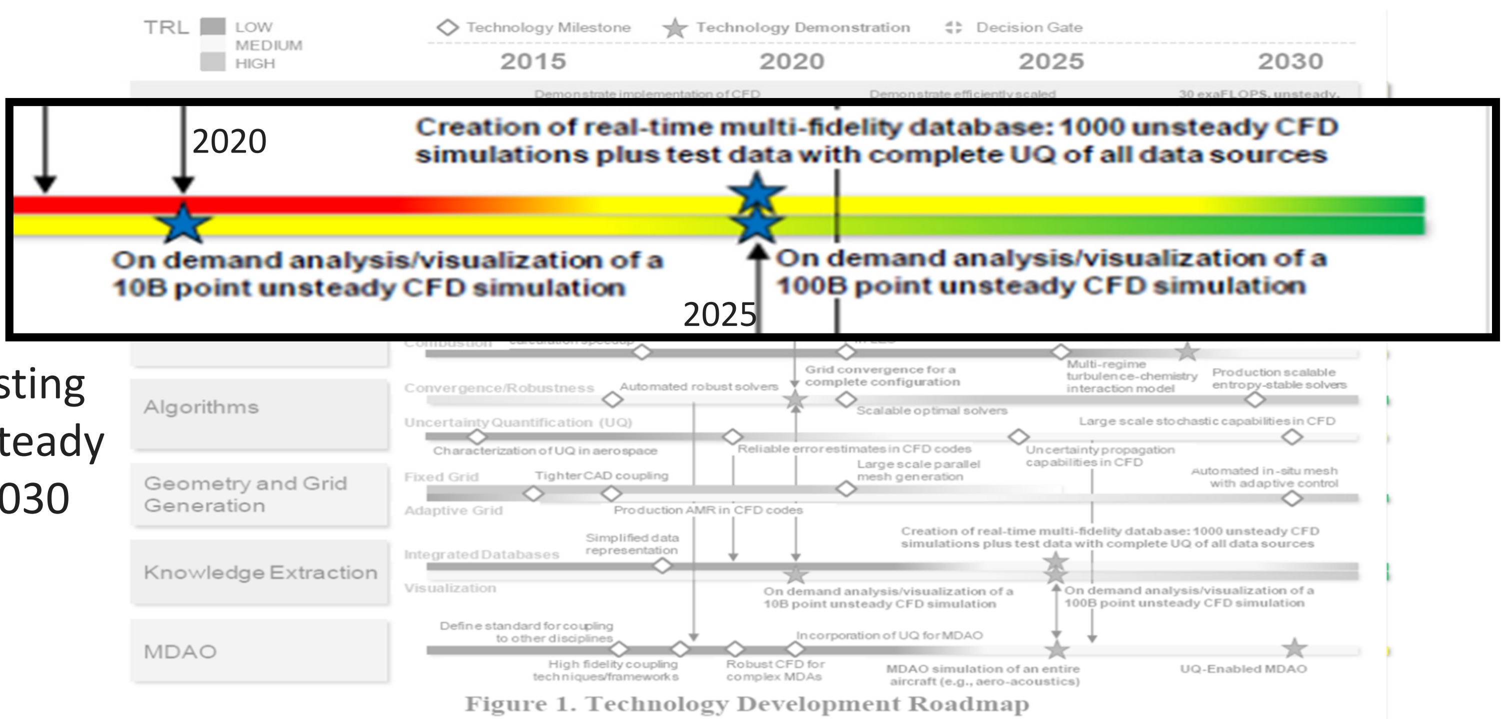

Figure 1. Technology Development Roadmap

NASA is forecasting trillion cell unsteady CFD cases by 2030

SZL Technology:

Pipeline

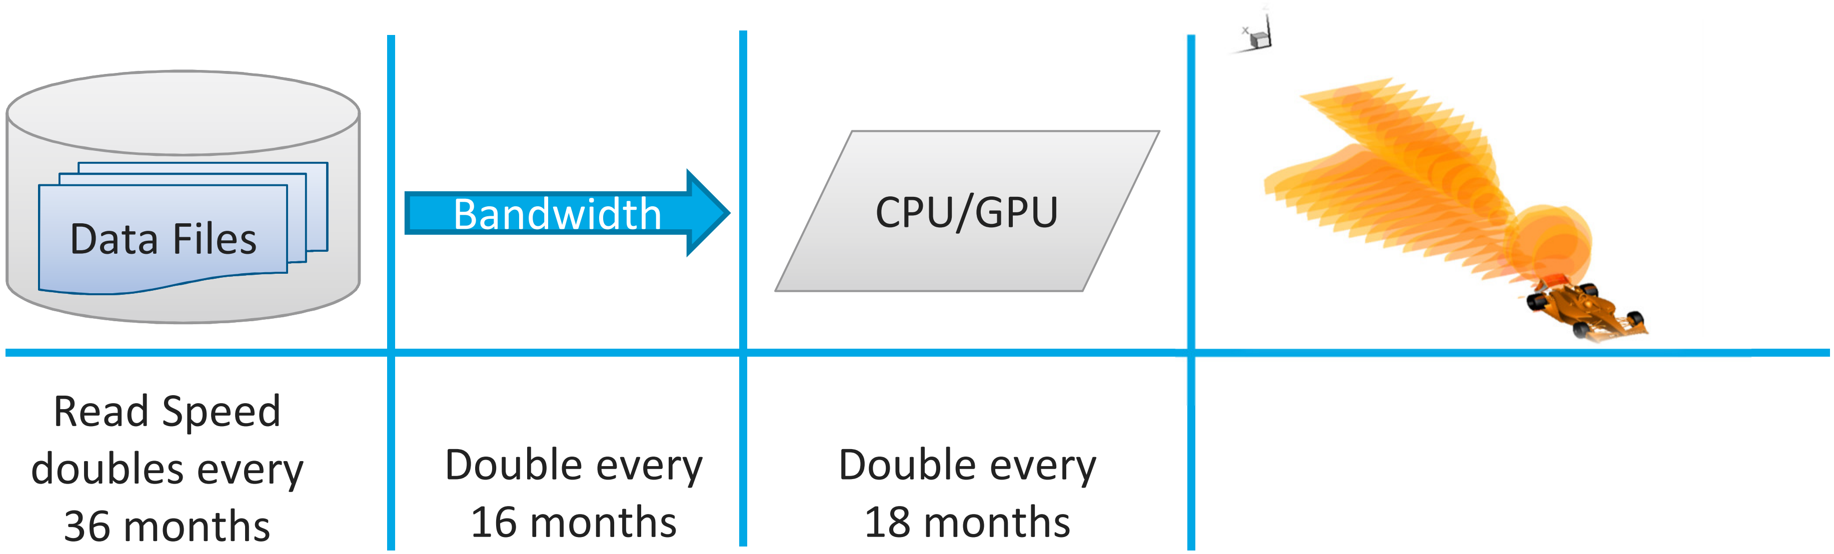

Data IO is the rate-determining step in the visualization pipeline.

SZL Technology:

Impact of Disk I/O Bottleneck

- Disk read performance growing slower than grid size
- Current visualization architectures will perform dramatically worse as time goes on!

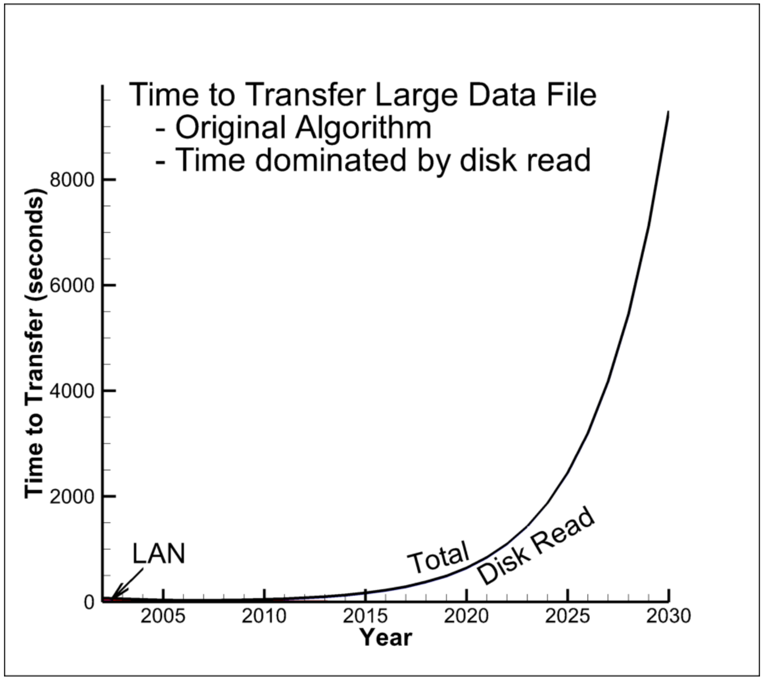

SZL Technology:

Basic Idea

- Reduce the amount of data you read
  - Must scale sub-linearly with the size of the grid
- Subzone Load-on-Demand (SZL)
  - Save indexed data file
  - Load only the data you need (Lazy Loading)

SZL Technology:

How SZL Works

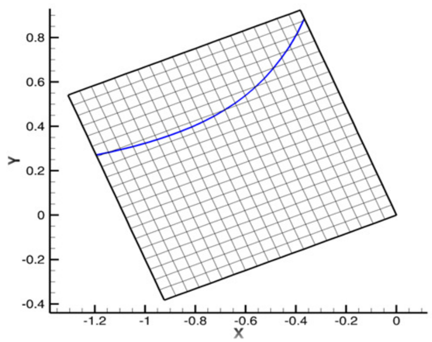

Example 2D contour line

- Current methods require loading data for zone
- For large data loading can be time intensive

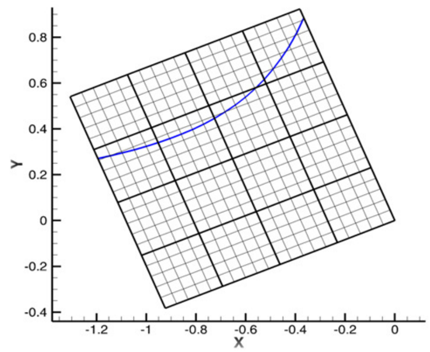

Domain can be indexed

- Decomposition of domain into smaller subdomains
- These subdomains can be indexed

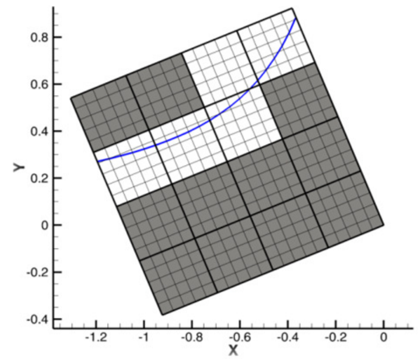

Data required for line is 5/16 of total data

- Loading time reduced
- Memory requirements reduced

SZL Technology:

Performance Improvements

- Amount of data used is dramatically less with SZL
  - SZL advantage grows exponentially with time
- Transfer time grows much slower than before

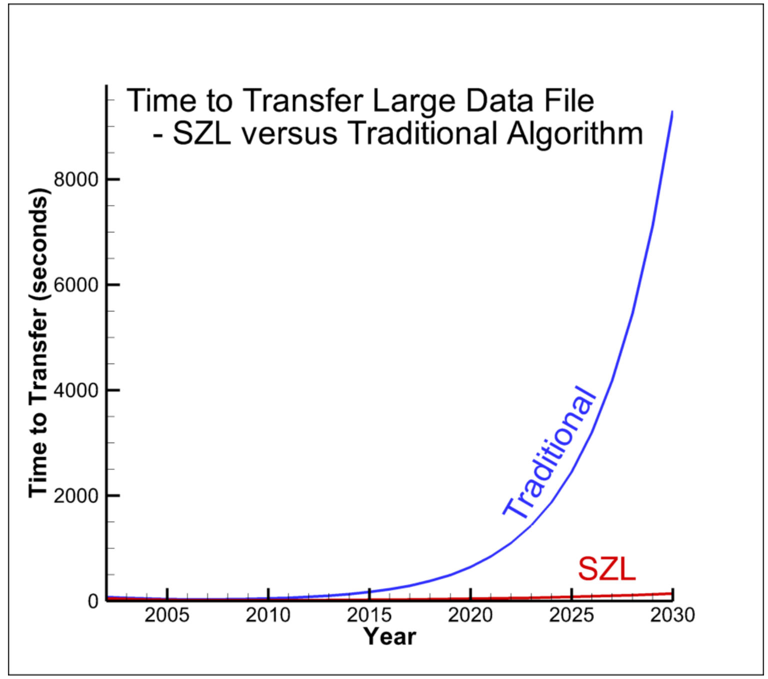

SZL Technology:

Performance Improvements

- 100% of the test show improved performance with SZL

- Results vary depending on specific data set, but all but the smallest test showed >10x speedups (FE data)

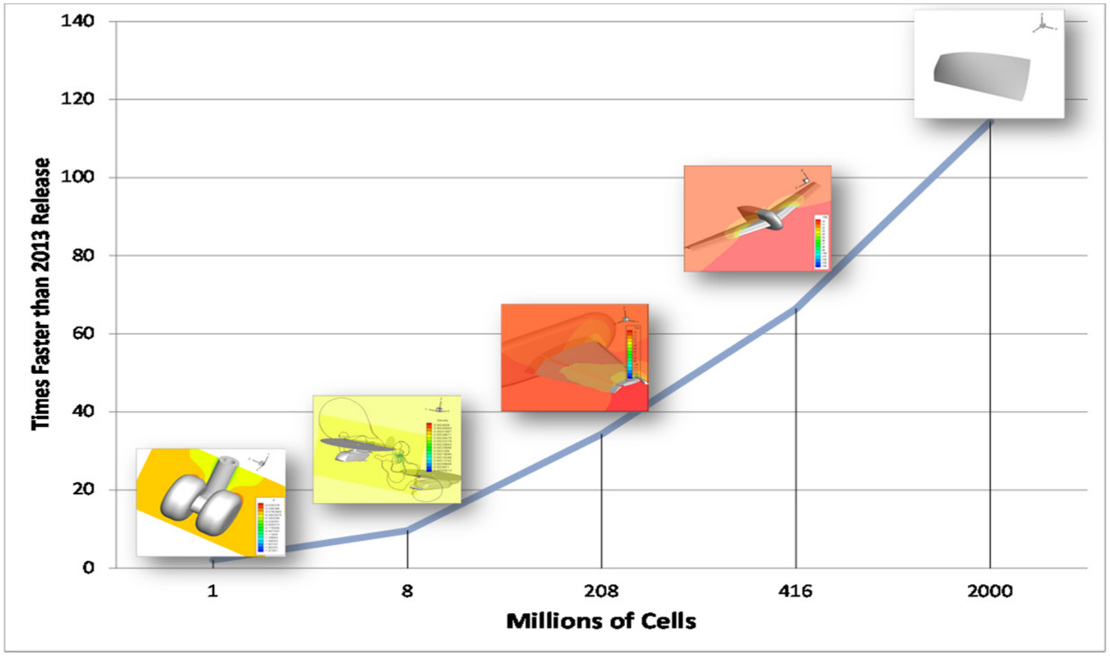

SZL Technology:

Memory Reduction

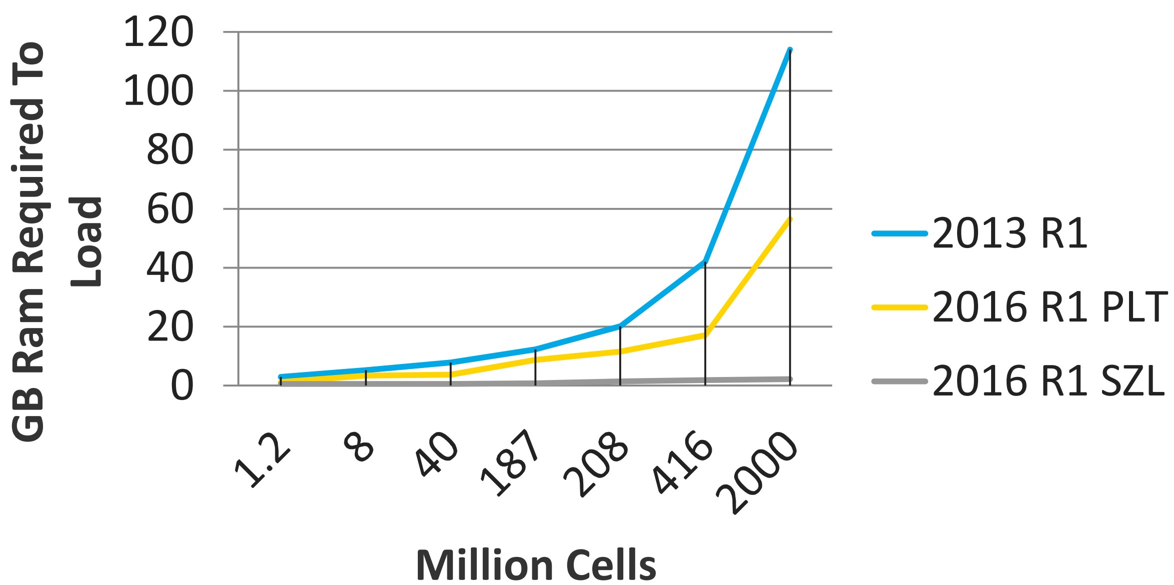

Memory reduction is what makes the impossible possible. With SZL you can load extremely large solutions on machines with limited memory.

SZL Technology:

Trillion Cell Challenge

- Achieved our goal of visualizing one-trillion cells in 2015
- Scaling is as predicted
- *Note: it took 3.5 days to write the 8.5TB file*

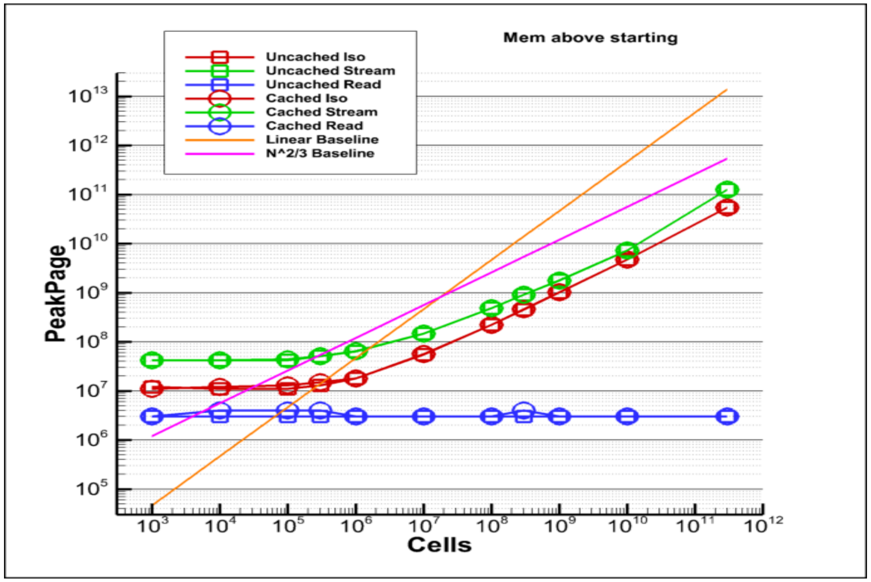

SZL Technology:

In-Situ Visualization/Post-Processing

- In-Situ: (post-processing) performed in place (in the CFD code)
  - Sometimes the post-processing is done on other nodes of the HPC system
- Desired Results:
  - Dramatically reduce the size of data written to disk (overcome the I/O bottleneck when writing).
  - Minimally impact the performance of the CFD code (CPU cycles and memory)
  - Maximize “explore-ability” of In Situ output file

SZL Technology:

Traditional In-Situ

- Extract features such as slices and iso-surfaces as surface data and save to file
  - Requires significant resources (CPU cycles and code space) to extract features
  - Sometimes the surface data is rendered to create an image (more processing, less data written)
  - If you made a mistake, must rerun (restart) CFD code

SZL Technology:

Subzone Based In-Situ

- Write out only the volume subzones needed for desired features
- Select subzones you write based on a query
- Example queries:
  - Q-Criteria for Overset grids: (Q>0.0 AND IBlank = 1)
  - Cp isosurface and an x-slice: (Cp = -1 OR X=100)

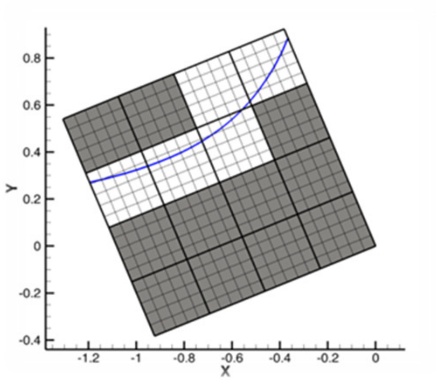

SZL Technology:

Comparison of Subzone Based and Traditional In-Situ

Desired Characteristics	Subzone-Based	Traditional
Minimize File Size	Larger (3x to 4x)	Smaller
Minimal HPC processing	Less	More
Post-write Exploration	Yes (some)	Less (view changes)
Post-write field derivatives	Yes	No

SZL Technology:

Results: Wind-Turbine LES

- OVERFLOW 2.2
  - OVERFLOW-D mode
  - Domain Connectivity Function (DCF)
  - Geometry Manipulation Protocol (GMP)
  - 16 near-body blocks (2.6M points)
  - Adaptive Mesh Refinement
  - 2nd-order differencing near-body
  - 4th-order differencing off-body
- Results
  - 10 m/s aligned with turbine
  - time step = 7152, 260M nodes in 5600 blocks

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Q=0.0
- Too much

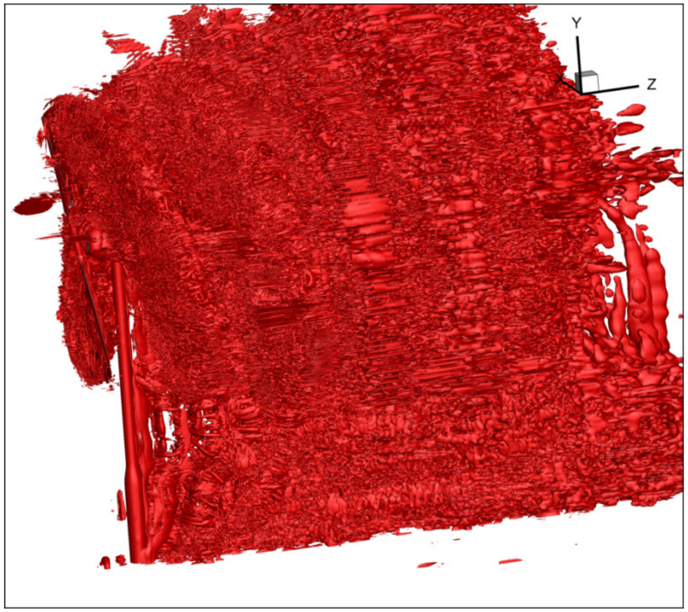

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Q=0.001
- Better

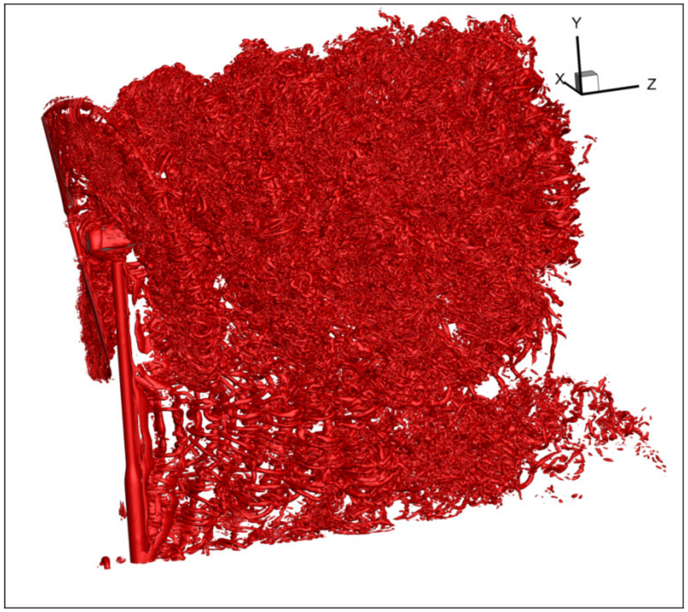

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Q=0.003
- Good

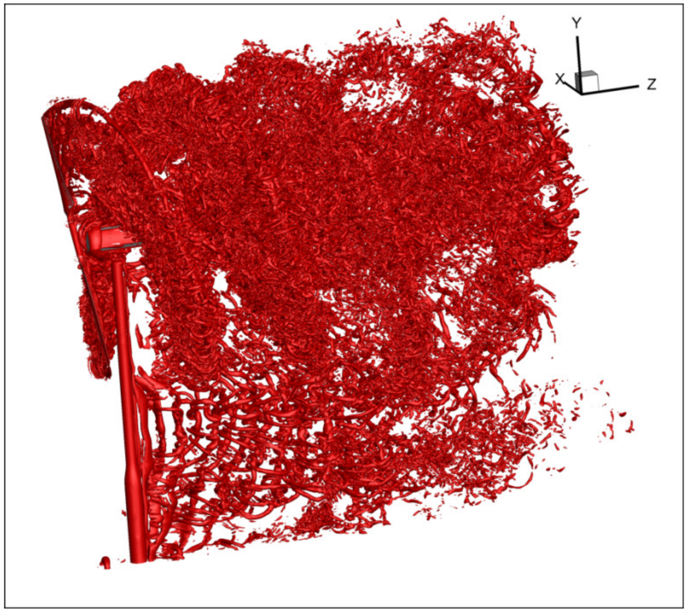

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Q=0.01
- Good (sparse)

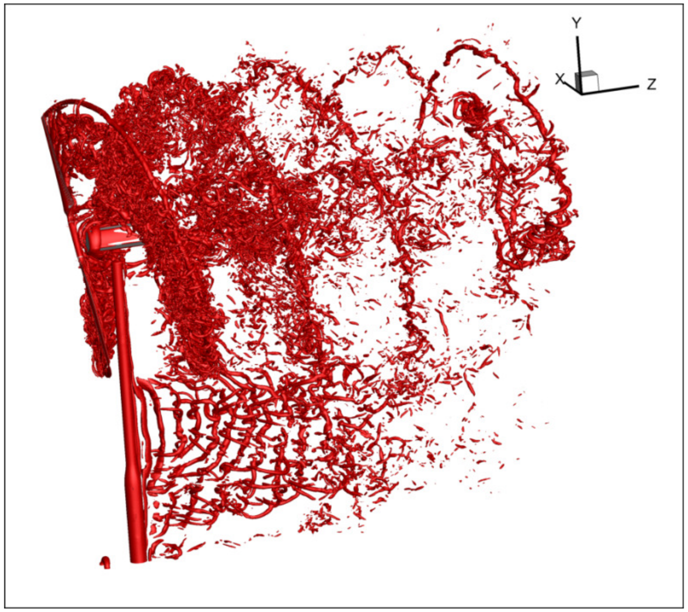

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Q=0.03
- Too sparse

SZL Technology:

Need for Exploration

- Example: Q-Criteria for LES of wind-turbine wake
- Rather than creating five separate In Situ Iso-surface extractions, with Subzone Based In Situ you can
  - Specify a range of Q to include in the file (say 0.003 to 0.03)
  - Interactively explore the data to find the best value of Q.
    - Probably between 0.003 and 0.01

SZL Technology:

Results: Wind-Turbine LES

- File Size (chart)
  - SZL In Situ files are 3x to 5x larger than traditional files
  - SZL In Situ files are much smaller than full SZL file
- Iso-surface details
  - For Q=0.01, 6.6M triangles and 3.4M points
- Load Times and Memory (For Q=0.01)
  - SZL: 70.8 sec and 5.2 GB
  - SZL In Situ: 11.8 sec and 3.5 GB

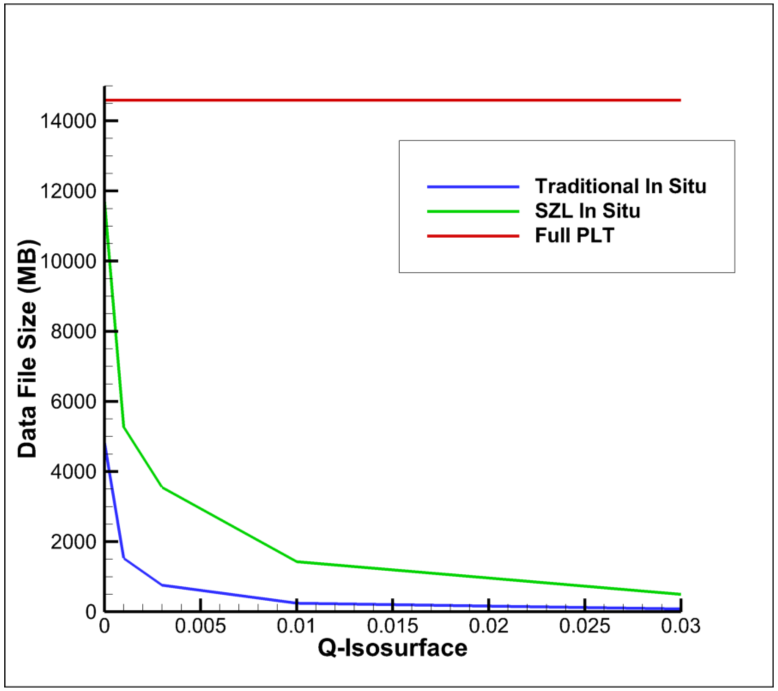

SZL Technology:

Results: NASA Trapezoidal Wing

- Geometry (High-Lift Prediction Workshop)
  - Trapezoidal wing with flaps and slats extended
  - Half body and mirrored to create full body
- Grid
  - Finite-element (Bricks,Tets, Prisms)
  - Half: 204 Million Cells
- Solution
  - File - 12.9GB

SZL Technology:

Results: NASA Trap Wing Iso-surface

- File Sizes
  - PLT: 9.48 GB
  - SZL: 5.92 GB (node-map compression)
  - SZL In Situ: 0.64 GB (10.8% of full SZL)
  - Just Isosurf: 0.083 GB
- Load Times & Memory
  - PLT: 98 sec and 16.5 GB memory
  - SZL: 30.3 sec and 2.4 GB memory
  - SZL In Situ: 9.4 sec and 1.9 GB memory

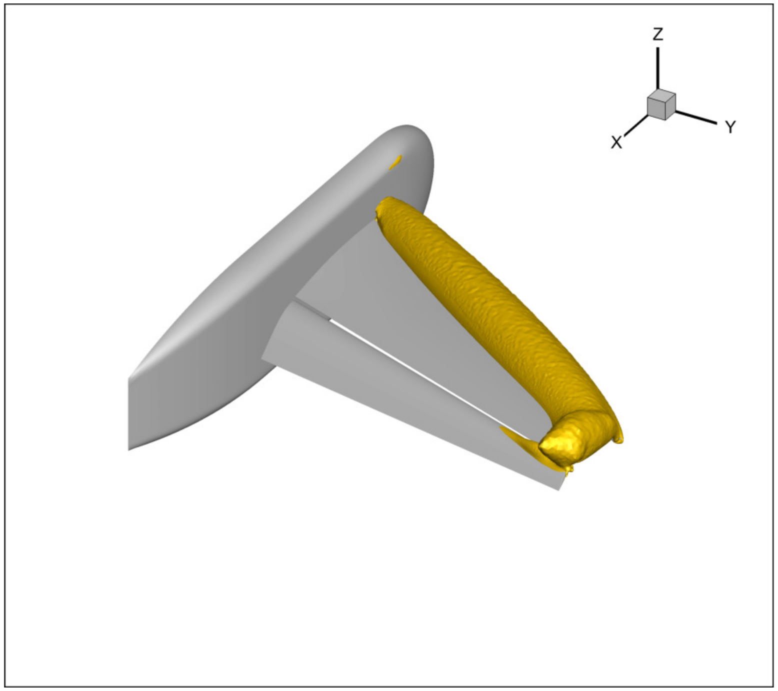

SZL Technology:

Results: NASA Trap Wing Slice

- File Sizes
  - PLT: 9.48 GB
  - SZL: 5.92 GB (node-map compression)
  - SZL In Situ: 0.075 GB (1.3% of full SZL)
  - Just Isosurf: 0.008 GB
- Load Times & Memory
  - PLT: 97 sec and 16.5 GB memory
  - SZL: 8.5 sec and 0.47 GB memory
  - SZL In Situ: 1.6 sec and 0.4 GB memory

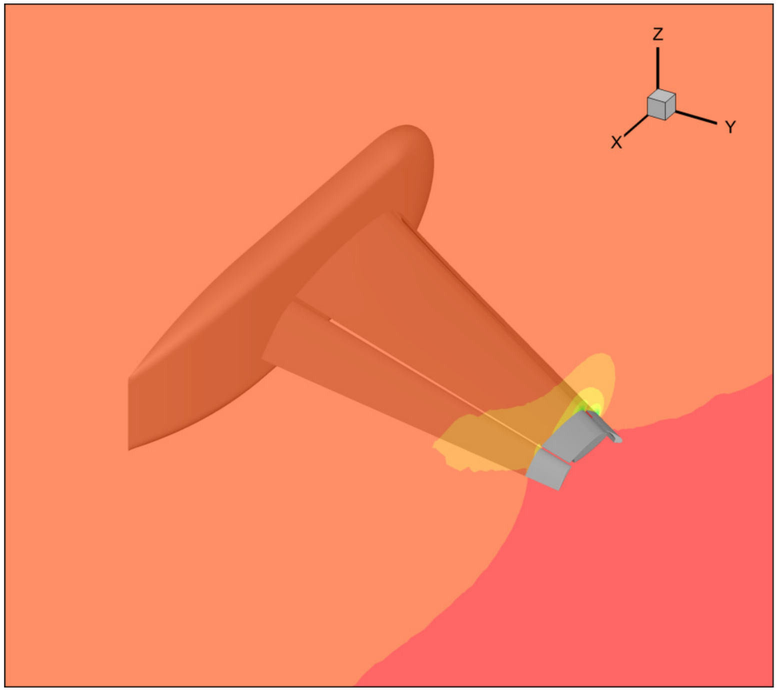

Summary:

Subzone-Based In-Situ:

- File Size
  - Much smaller than full SZL file
  - Bigger than traditional In-Situ
- Requires little processing on the HPC
- Allows post-write explorations and processing (derivatives)
- Loading Subzone In-Situ files into Tecplot
  - Faster and uses less memory than SZL
  - Much Faster and uses much less memory than PLT

SZL Technology:

Subzone-Based Client-Server

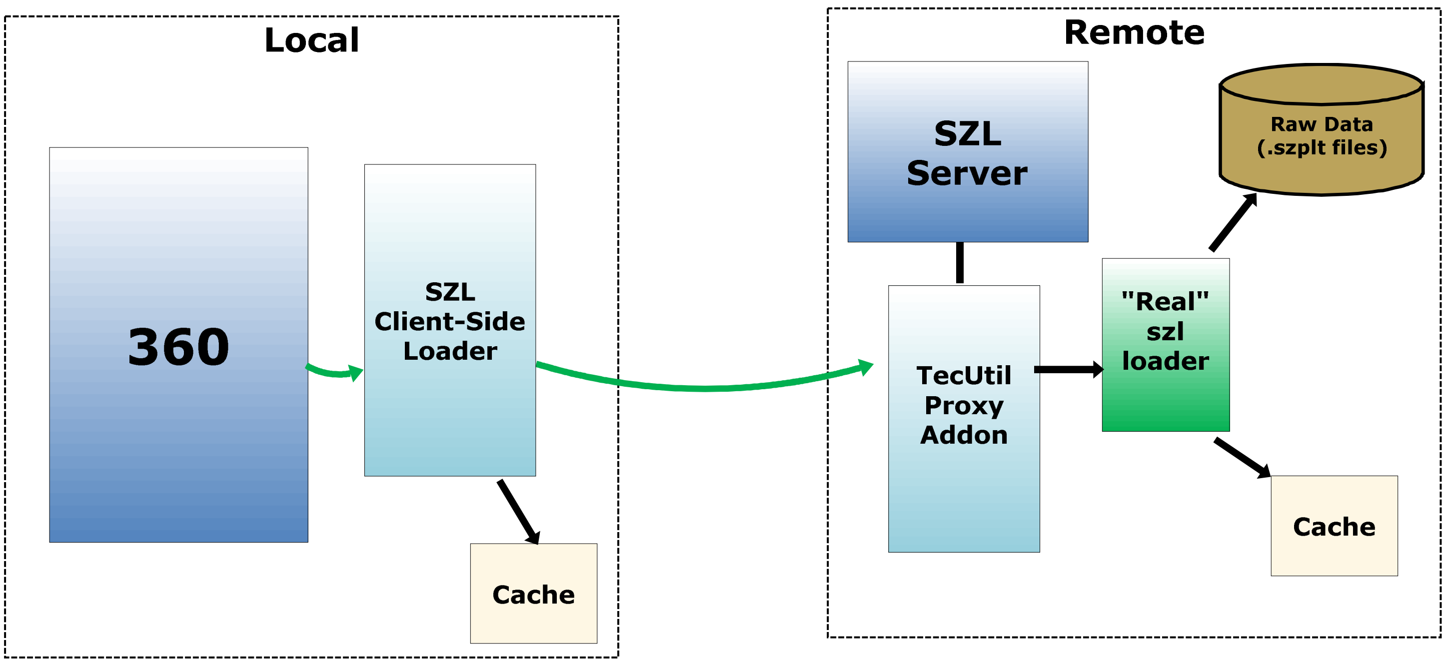

- Serves subzones instead of surfaces

SZL Technology:

Future Development

- In-Situ is currently proof of concept, need to complete
  - Add In-Situ queries to TecIO
  - Optimize write performance