



Turbomachinery R&D Acceleration using Titan

R. Srinivasan

Aero/Thermodynamic Engineer

Dresser-Rand

GTC 2015, San Jose, CA

DRESSER-RAND

Outline

- Background
  - Software Developments
  - Application Strategy
- Titan focused development
- Applications
  - Open Science
  - Turbomachinery R&D
- Conclusions

DRESSER-RAND

BACKGROUND

DRESSER-RAND

Background

- Partnership with Oak Ridge National Laboratory's Leadership Computing Facility
- Utilize large scale DoE hardware at OLCF to improve development time of new turbomachinery designs
- Commercial software package
  - The Fine/Turbo suite
    - CFD tools for high fidelity turbomachinery analysis
    - Advanced features for fast, high fidelity simulations
- Software development program
  - Jointly funded
- Time-allocation through DD and ALCC

DRESSER-RAND

Background – Software Developments

- Virtual decomposition of the multi-block grid
- Distributed memory parallel computing with MPI
- I/O improvements
- Parallelization of turbomachinery related BC
- Efficient scalability with up to 4000 compute processes
- Ongoing effort

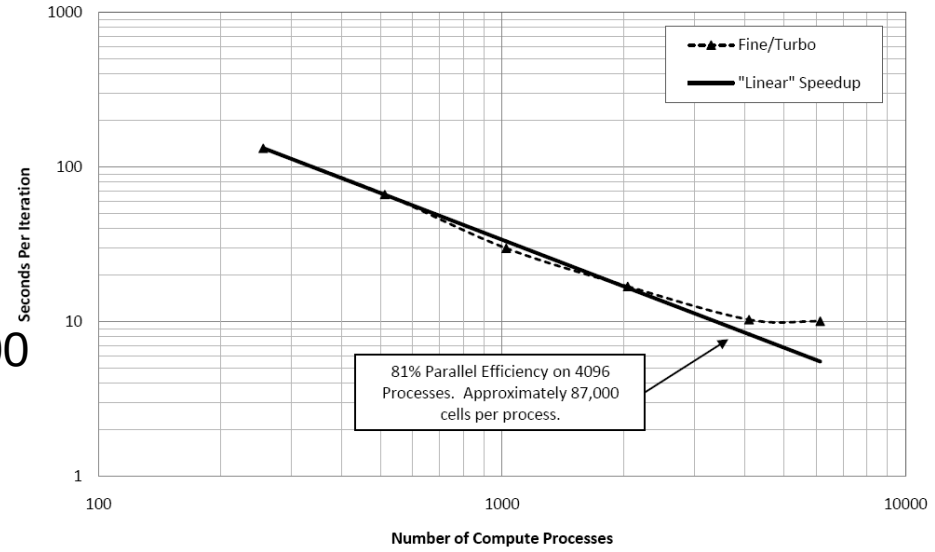

DRESSER-RAND

Application Strategy

- Turbomachinery R&D
- Time frame for design and development
- Parametric optimization
- Ensemble runs for database generation
- Design refinement

DRESSER-RAND

TITAN FOCUSSED DEVELOPMENT

DRESSER-RAND

GPU Strategy

- Utilization of new GPU hardware on Titan
- Software Limitations
- Industrial Limitations
- Cost of rewriting the solver
- Investigated alternatives
  - Selection of appropriate code module for acceleration

DRESSER-RAND

GPU Accelerated Module

- Advanced residual smoothing algorithm, bringing a level of implicitness to the time integration solver
- Permits the use of high CFL numbers
- Yields a drastic decrease in the number of iterations needed for convergence (~4x)
- Time per iteration is increased due to the additional module arithmetic (~2x)
- Attractive candidate for threaded execution, either on CPU or GPU

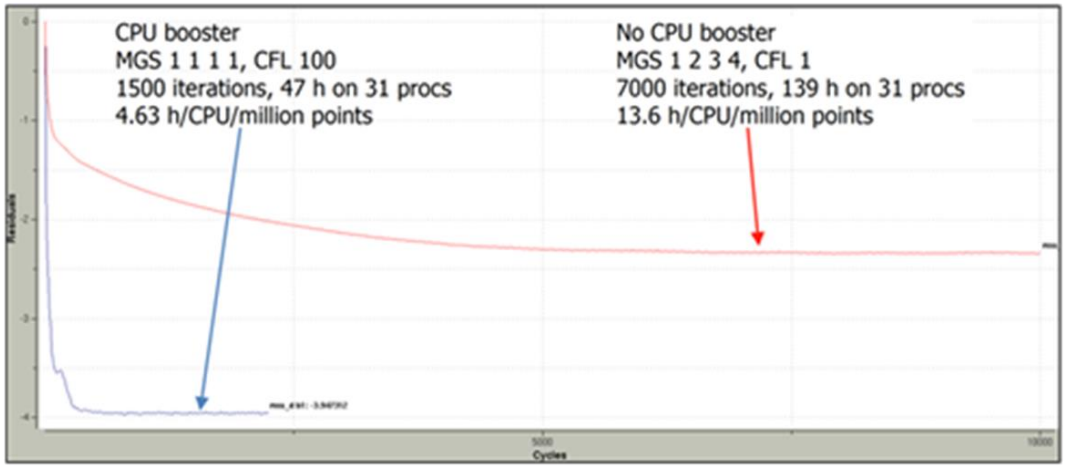

DRESSER-RAND

Hardware and Software Considerations

- Titan nodes are equipped with PCIe 2.0
- Each node contains 16 CPU cores, all of which will simultaneously transfer data to a single GPU
- Bandwidth in/out of the GPU is a bottleneck
- GPU memory is also a limitation for larger computations
- OpenACC programming model
  - Directive based approach for usability and ease of maintainance

DRESSER-RAND

Test Case Performance

- Positive speedup of 1.25x observed during execution of typical sample model
- Speedup dependence on internal boundary conditions
- Peak speedup of 1.7x observed for conformal BC model

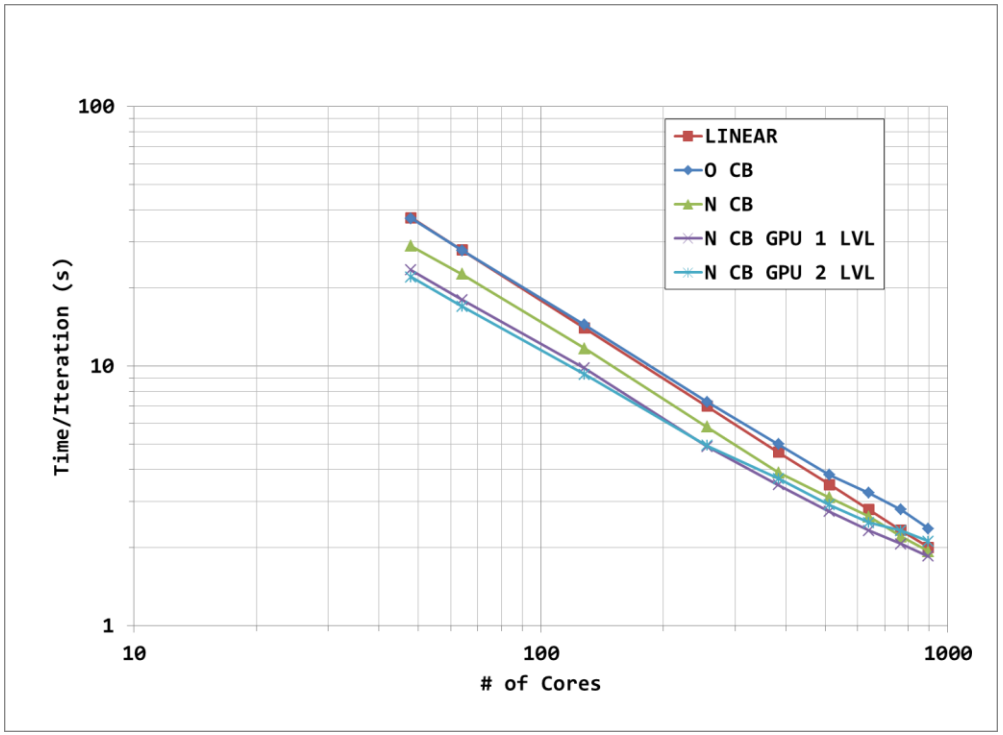

Typical sample speedup of 1.25x, maximum of 1.7x.

DRESSER-RAND

IO Performance

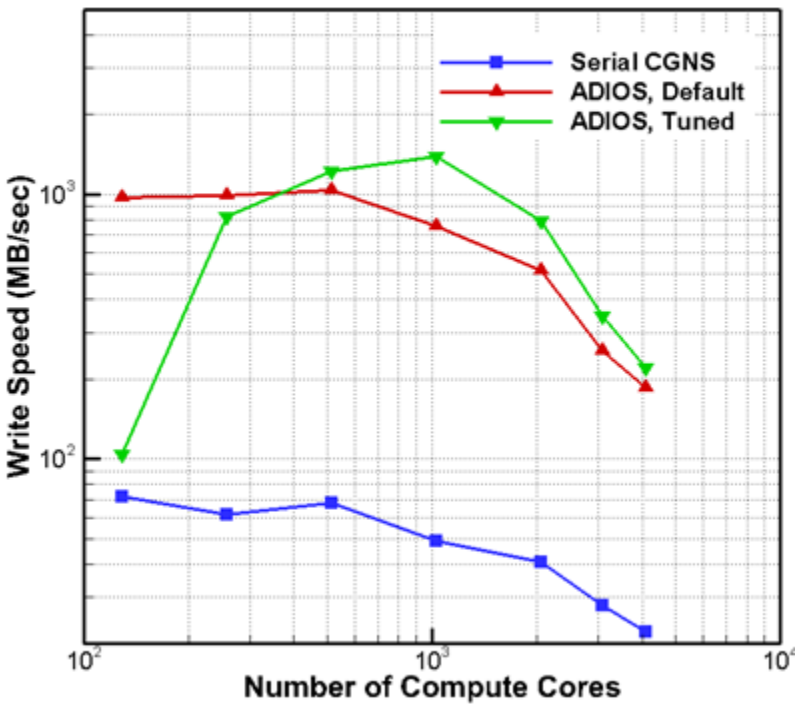

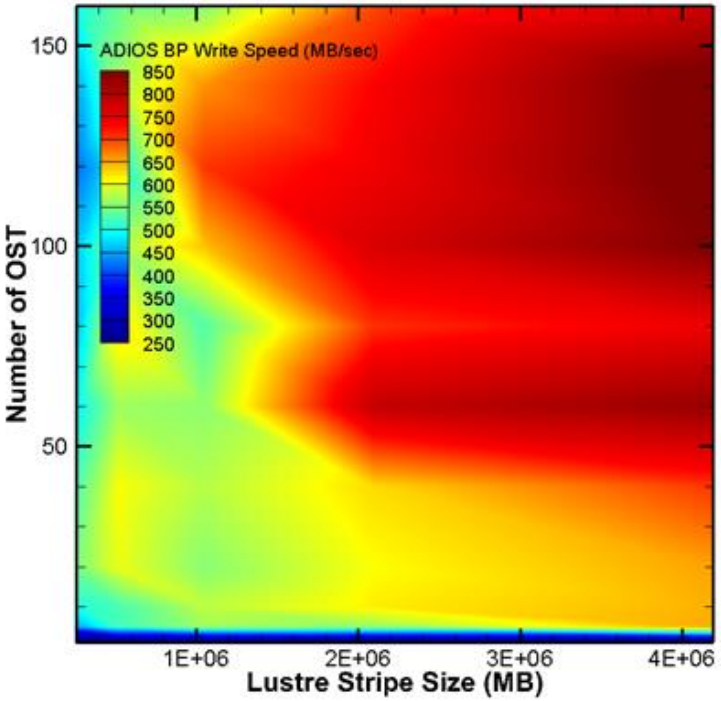

Order of magnitude increase in write speed.

DRESSER-RAND

APPLICATIONS

DRESSER-RAND

Open Science

- Supersonic inlet/isolator model
- Shockwave boundary layer interaction
- Conformal block boundaries
- Experimental comparison
- On-going work

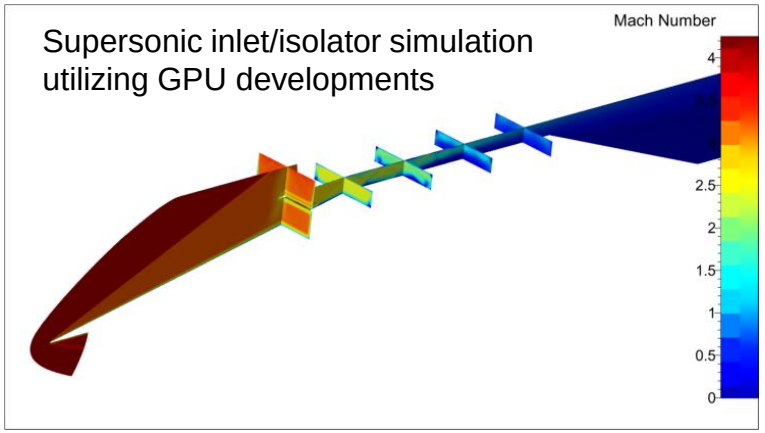

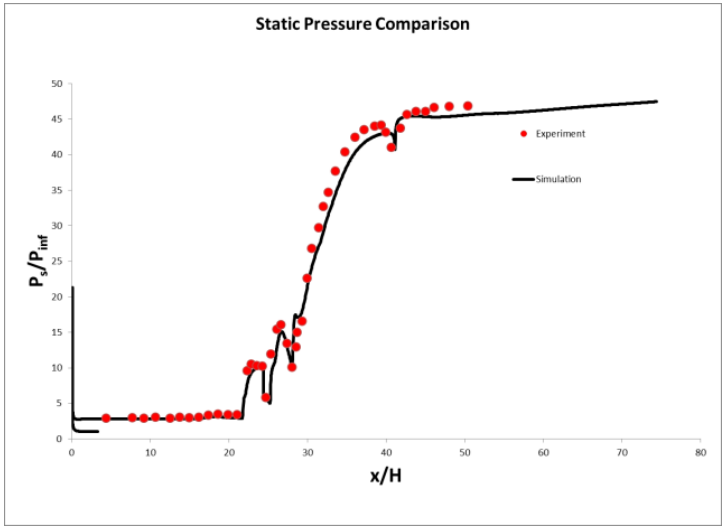

DRESSER-RAND®

Optimization

- D-R utilizes optimization techniques for enhancing turbomachinery design
- Results for perturbed database designs
- Analyze multi-dimensional data
- Database refinement
- Tolerance sensitivity

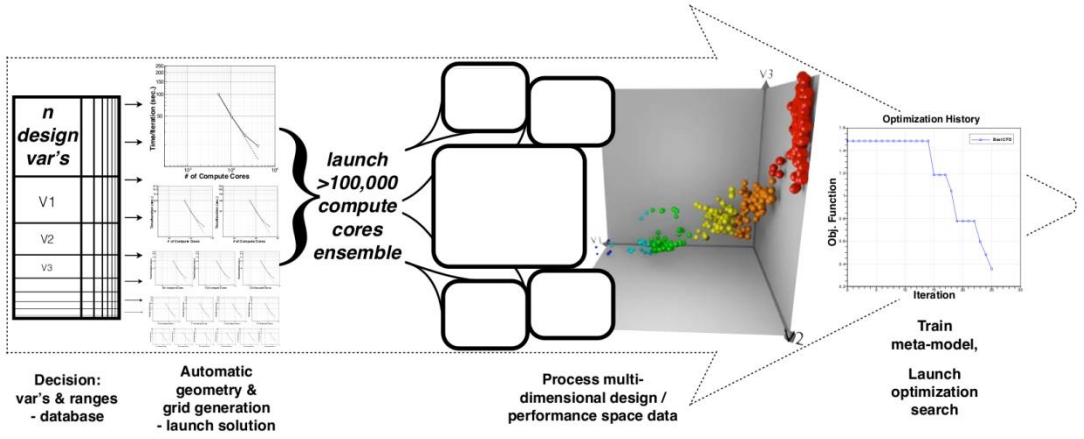

Titan – An enabling technology for accelerating turbomachinery R&D.

DRESSER-RAND®

Impact on Time-to-Solution

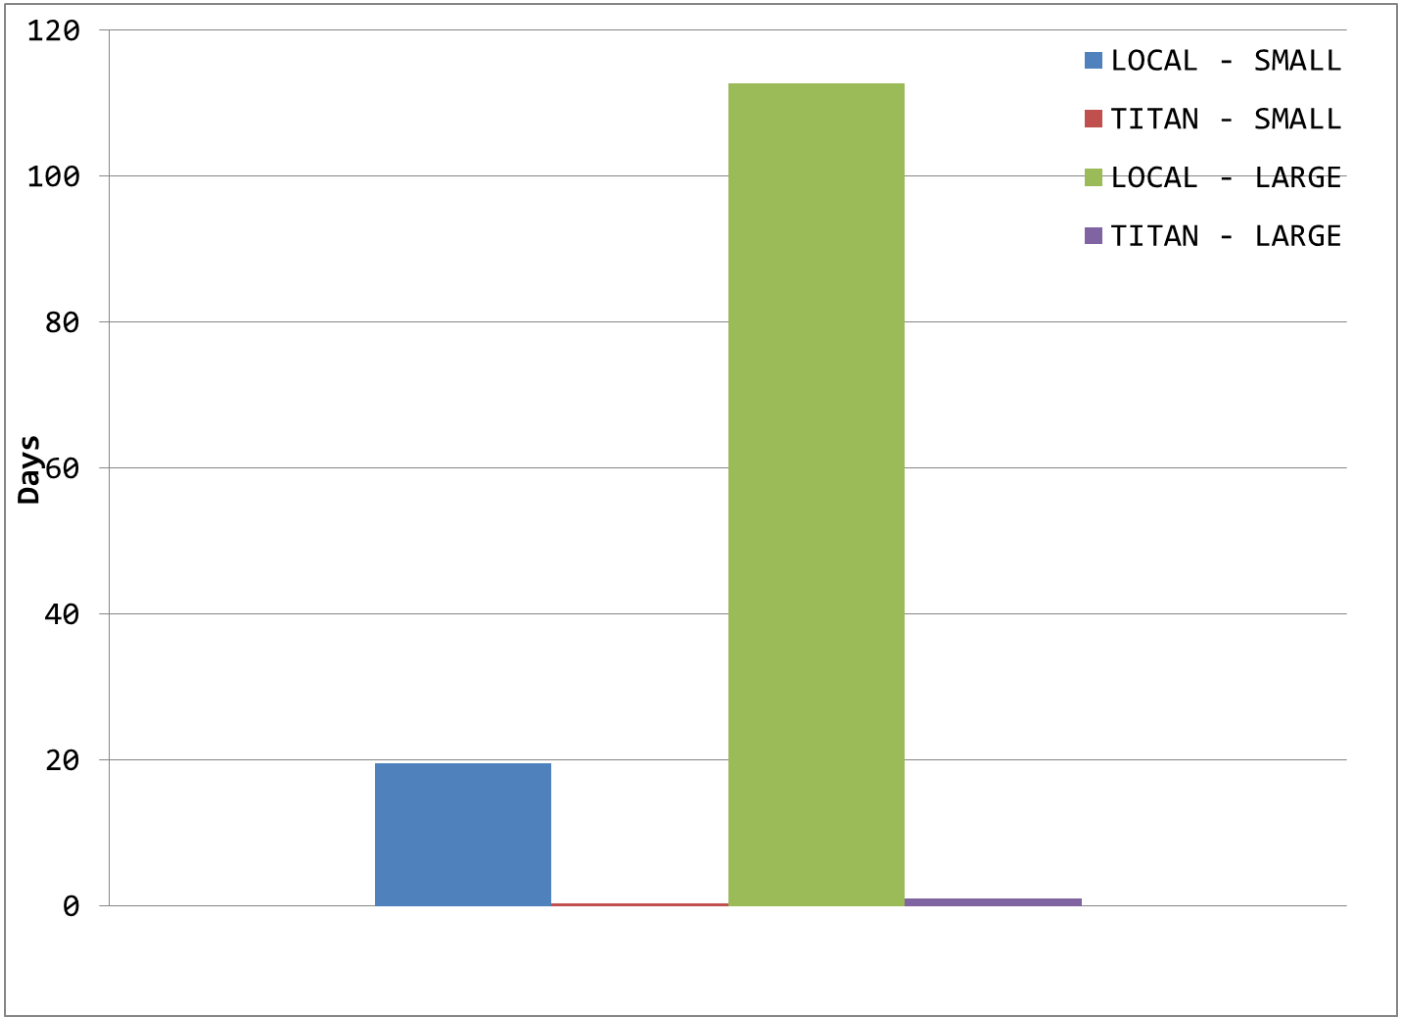

DRESSER-RAND

Summary

- Development efforts to enable scalable solver execution on OLCF hardware.
- The GPU acceleration of the solver module
  - The CPUBooster convergence acceleration module ported to GPU
  - Restructuring and instrumentation of the CPUBooster module with OpenACC directives has yielded a global iteration speedup of 1.2X-1.7X
  - With proper code restructuring, OpenACC allows easy GPU acceleration

DRESSER-RAND®

Summary

- The use of the CPUBooster module along with the Titan GPUs results in a 2.5X decrease in turnaround time and resource usage for D-R's optimization work.
- Large scale optimization related simulations are feasible due to availability of access to Titan
- Successful porting of the CPUBooster module has led to a new development program to port additional solver modules to execute on GPUs.

Acknowledgements

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DEAC05-00OR22725.

- Dr. Jack Wells, Director of Science, ORNL
- Suzy Tichenor, Director of Industrial Partnerships, ORNL
- David Gutzwiller and Mathieu Gontier, Numeca
- John Levesque (CRAY)
- Jeff Larkin (Nvidia)
- Dresser-Rand Employees

DRESSER-RAND





Questions

DRESSER-RAND