## Turbulent Flow Solvers – Perspectives on HPC and Solver Implementations

#### UMICHIGAN / NASA SYMPOSIUM ON ADVANCES IN TURBULENCE MODELING



ANN ARBOR, MI July 12, 2017

Juan J. Alonso aerospacedesignlab

Department of Aeronautics & Astronautics Stanford University







## Turbulent Flow CFD Applications Are Many...

 Computational methods to aid in engineering analysis and design are <u>pervasive</u>: higher importance in future

### • Applications are many:

- Gaining basic understanding of phenomena
- Understanding multi-disciplinary interactions
- Guiding design / optimization
- Certification by Analysis / reducing margins for design

### • Fundamental questions:

- How do we <u>embed high-fidelity</u> methods in every-day design processes?
- How do we <u>automatically</u> manage errors and uncertainties?
- How do we <u>leverage future computing power</u>?
- How do we show <u>industrial relevance</u>?
- How good is good enough? And when?

## ADL - Main Research Areas

### Multidisciplinary Analysis and Optimization (MDO)

- Managing tool fidelity / multi-fidelity approaches
- Design under uncertainty (robust / reliability-based design)
- Hierarchical decomposition methods
- Validation & Verification (V&V) and Uncertainty Quantification (UQ)
  - Management of numerical errors
  - Propagation of natural variability uncertainties
  - Understanding of model-form uncertainties
- System/Vehicle-Level Implications
  - Problems are not just at the component level
  - System-level interactions are fundamental
  - Stochastics, interactions, strategic players/actors

## Problems of Interest ... In Pictures











## CFD Vision 2030 Study

### Emphasis on physics-based, predictive modeling

Transition, turbulence, separation, unsteady/time-accurate, chemicallyreacting flows, radiation, heat transfer, acoustics and constitutive models

### Management of errors and uncertainties

Quantification of errors and uncertainties arising from physical models, mesh and discretization, and natural variability

### o Automation in all steps of the analysis process

Geometry creation, meshing, large databases of simulation results, extraction and understanding of the vast amounts of information

### Harness exascale HPC architectures

Multiple memory hierarchies, latencies, bandwidths, programming paradigms and runtime environments, etc.

### Seamless integration with multi-disciplinary analyses and optimizations

High fidelity CFD tools, interfaces, coupling approaches, the science of integration, etc.

Slotnik, et al., "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," NASA/CR-2014-218178, 2014







## CFD Vision 2030 Study

### Emphasis on physics-based, predictive modeling

Transition, turbulence, separation, unsteady/time-accurate, chemicallyreacting flows, radiation, heat transfer, acoustics and constitutive models

### o Management of errors and uncertainties

Quantification of errors and uncertainties arising from physical models, mesh and discretization, and natural variability

### o Automation in all steps of the analysis process

Geometry creation, meshing, large databases of simulation results, extraction and understanding of the vast amounts of information

### Harness exascale HPC architectures

Multiple memory hierarchies, latencies, bandwidths, programming paradigms and runtime environments, etc.

### Seamless integration with multi-disciplinary analyses and optimizations

High fidelity CFD tools, interfaces, coupling approaches, the science of integration, etc.

Slotnik, et al., "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," NASA/CR-2014-218178, 2014







## Predictive Computational Science: V&V and UQ



Quantifying discretization errors is a first step to quantify sources of uncertainty. Understanding uncertainties is necessary to achieve certification. Stanford University

# Physical Modeling Issues in Future CFD Solvers

### RANS model based:

- Robustness improvements
- Speed of convergence (steady <u>and</u> unsteady RANS)
- Overall simulation cost
- Complex geometry representation
- Sensitivities

### Scale-resolving methods:

- Low-dissipation spatial discretizations
- Accurate time-stepping methods
- Subgrid-scale and turbulence models, wall models
- Sensitivities of unsteady flows!

### • Data-driven solvers:

- All of the above plus...
- Rapid embedded querying of "learned" portions of the model
- Advanced model data decomposition techniques (for very large databases)
- Offline: scalable machine learning tools common to other communities

## Compute Hierarchy in Hardware









#### SPECIFICATIONS

GPU Architecture	NVIDIA Pascal
NVIDIA CUDA® Cores	3584
Double-Precision Performance	4.7 TeraFLOPS
Single-Precision Performance	9.3 TeraFLOPS
Half-Precision Performance	18.7 TeraFLOPS





## Memory Hierarchy in Hardware



## The Compute-Memory Gap



\*Source: extremetech.com

## The Performance Equation...



- Must tackle all of these elements in order to obtain scalable and performant code
- A non-trivial effort ... that nobody is interested in! Stanford University

## But Progress <u>IS</u> Slow...

- These are long-term research issues...
- ... but it still takes a significant amount of time before we can transition turbulence research/HPC results to industrial applications
- We have been doing an "experiment":



- SU2: Analysis and Design Optimization in Complex Configurations
- Are <u>community codes</u> well placed to accelerate the transition of turbulence research and HPC implementations?

### What is SU2?



The Open-Source CFD Code

The SU2 suite is an **open-source** collection of C++ based software for multi-physics simulation and design on unstructured meshes (i.e., CFD!).

First and foremost: a Community code!

SU2 is under active development at Stanford University in the Department of Aeronautics and Astronautics and **in many places around the world**.

http://su2.stanford.edu https://github.com/su2code/SU2

2015 SU2 Team, "SU2: An Open-Source Suite for Multi-Physics Simulation and Design," AIAA Journal, 2015, doi: 10.2514/1.J053813.
 2012 SU<sup>2</sup> team, "Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design", AIAA Paper 2013-0287.
 2013 SU<sup>2</sup> team, "Stanford University Unstructured (SU2): Open-source analysis and design technology for turbulent flows", AIAA Paper 2014-0243.

#### SU2 – Multi-Physics Analysis and Design Since release in Jan. 2012: 500,000+ web visits across 178 countries 10s of thousands of downloads, 15,000+ email addresses on the user list 200+ forks, 170+ developers in list Scientific Computing TE TECHNISC TU-P&P: Non-Ideal Compressible Fluid-**Dynamics** (NICFD) NICFD → branch of Fluid-Mechanics studying the flow physics of upercritical flows, dense vapors, and two-phase flows Imperial College London Partitioned FSI problem DDES + FWH Fluid Structure $p_{\Gamma}, \tau_{\Gamma}$ $\rho_f, \mathbf{v}, E$ u, ů Interface TUDelft $\mathbf{u}_{\Gamma}, \dot{\mathbf{u}}_{\Gamma}$ $p_{\Gamma}, \bar{\bar{\tau}}_{\Gamma}, \mathbf{u}_{\Gamma}, \dot{\mathbf{u}}_{\Gamma}$ Mesh $\mathbf{u}_{\Omega}, \dot{\mathbf{u}}_{\Omega}$ **DG-FEM Higher-Order Solver**



Lines of Code in SU2 by Release (w/out comments or blanks)

\*includes code in externals/

## Now, It Is Not All About Research...HPC!

• Based on several recent experiences:

- Intel Parallel Computing Center (IPCC) at Stanford
- Argonne National Lab, Theta Early Science Program (ESP)
- Argonne National Lab, Aurora Early Science Program (ESP)
- Can community codes also accelerate transition of optimized HPC implementations?
- The effort (and knowledge) required to obtain scalable and performant CFD codes has increased substantially:
  - Large teams with varied expertise now required
  - Expertise rarely resides at Universities any more

## A Tale of Two Solvers In SU2...

- Workhorse of SU2 is a the 2<sup>nd</sup> order FV solver: SU2 FV
  - Unstructured mesh, edge-based
  - Median-dual control volumes
  - Many flux discretizations and reconstructions
  - RANS (SA/SST, DDES variants)
- For the past 2 years, we have been developing a DG-FEM solver based on the "same" infrastructure: SU2 DG-FEM
  - Element-wise data structure, arbitrary order
  - Flux computations at faces
  - Higher arithmetic intensity
  - feature\_hom, to be released this summer
  - Targeted to LES and WMLES
- Multi-year efforts, jointly with Intel, to improve performance and scalability on Intel Xeon Phi (KNL) platforms





Stanford University



 $\partial \Omega$ 

Primal Grid

## Compute Patterns in SU2 FV

Compute Pattern	Challenges	Optimizations			
Edge-based Loops: capture the actual physics of the problem (residuals & Jacobians computation)	Partitioning, retaining spatial locality Loop-carried dependencies Irregular - mem accesses, working-sets, write contention	for all Edges do vertices forming the edge $-[v_1, v_2]$ $y[v1] + = \mathbf{F}(V[v1], V[v2])$ $y[v2] + = \mathbf{F}(V[v2], V[[v1]))$ end for			
Sparse recurrences: sparse linear algebra kernels (Solution Ax = b)	Limited parallelism, extracting concurrency Large load imbalances Low arithmetic intensity, BW- bound Irregular - mem accesses, working-sets	1 2 3 4 5 6 7 8 (a) Non-zero pattern of a lower triangular sparse matrix (b) Corresponding dependency graph of forward solve			
Collectives & Sync	Algorithm characteristics, Network b	ound			
Vertex-based loops	Vertex-wise concurrency, Low compute intensity, Bandwidth bound				
		Stanford University			

## Performance Optimization on Modern Platforms



## **SU2 Performance Optimization Test Cases**



We are now routinely running large-scale parallel computations with 5000+ ranks for numerical performance experiments

## Key Lessons Learned – SU2 FV Solver

- Edge-loops/Face-loops and Sparse linear algebra bulk of execution time
  - Primary overheads: Irregular access, limited parallelism, loadimbalance
- Imperative to exploit *fine-scale concurrency* among cores & *SIMD* 
  - Sub-domain decomposition vs. data decomposition
  - Avoid synchronizations (atomics, barriers...)
  - SIMD: Unit-strided memory access, outer-loop vectorization

### Memory optimizations

- Most kernels in FV CFD codes are memory BW bound
- Memory efficient algorithms RCM reordering
- AoS-to-SoA transformations Compact working sets
- Prefetching

### Performance Modelling – How high can we go?

### Amdahls's Law for Multicores



### Amdahls's Law for Vector Multicores

$$Speedup = \left(\frac{1}{Serial_{frac} + \frac{1 - Serial_{frac}}{NumCores}}\right) * \left(\frac{1}{Scalar_{frac} + \frac{1 - Scalar_{frac}}{VectorLength}}\right)$$

Ideal Speedup: NumCores\*VectorLength (requires zero scalar, zero serial work)

## Performance Modelling – Roofline Model







## Roofline Model – Crossover Point



 $CI = 2N^{3}/(24N^{2}) = N/12$ 

## SU2 Performance Improvements through:

- Using the cores
- Using SIMD
- Efficiently using the *memory*

## SU2 FV Performance Improvements



## Key Optimizations performed:

### 1. Utilizing the Cores fully:

- Hybridized the code with MPI+OpenMP
- High-level OpenMP

### 2. Efficient vectorization:

- Reduced gathersscatters in SIMD loops
- Outer-loop vectorization

## 3. Efficiently using the *memory*:

- AoS-to-SoA layout change
- RCM re-ordering

## **Timeline of Performance Improvements**



Major optimization milestones:

(a) + Vectorized Viscous\_Residual kernel, changed layout of Gradient\_Primitive to be more cache/simd friendly, hard-coded some loop trip counts, other general optimizations

(b) + Added routine for edge position vectors (instead of getting Coord\_i and Coord\_i), several other general optimizations

(c) + Manual fusing of viscous and convective residual edge loops, data and compute layout changes in residual kernels

(d) + Optimized MPI (using MPI\_Wait instead of MPI\_Waitall); added nowait (OpenMP) and nontemporal stores (SIMD) clauses

## Extracting the Parallelism from the Cores

### OpenMP threading – Example from SU2 CFD code



# Fine-Grained ||ism: Data Vs. Domain Decomposition

- Data Decomposition
- Option 1: <u>Basic partitioning with</u> <u>atomics</u> – Just use "#pragma omp parallel for" and use "#pragma omp atomic" when writing out; Can use dynamic scheduling
- Option 2: <u>Coloring or Level</u> <u>Scheduling</u> – "Color" edges which can be operated on concurrently (no two adjacent edges have the same color). Do multiple passes. (No need for atomics)



Edge coloring, Source: Wikipedia

# Fine-Grained ||ism: Data Vs. Domain Decomposition

Domain Decomposition – Using METIS Library

 <u>Break graph at edges</u> – Redundant compute on interfacial edges. *No atomics*. Owner thread of the boundary point updates the value.
 [Cons: Cannot do dynamic thread scheduling]



METIS decomposition







## Efficiently Using the *Memory*

AoS (array-of-structures) to SoA (structures-of-arrays) – Improves cache hits, enables contiguous memory access

```
class CVariable {
                                      double *p;
    double p;
                                      double *u, *v, *w;
    double u,v,w;
  }
 CVariable **node;
for (iEdge = 0; iEdge < Nedges; ++iEdge)</pre>
                                                for (iEdge = 0; iEdge < Nedges; ++iEdge)</pre>
{
                                                1
   iPoint = GetNode(0);
                                                   iPoint = GetNode(0);
                                                   Pres i = p[iPoint];
   Pres i = node[iPoint]->p;
   Uvel_i = node[iPoint]->u;
                                                   Uvel i = u[iPoint];
}
                                                }
```

## Efficiently using the *Memory*

Reverse Cuthill-McKee (RCM) re-ordering – Improves cache hits

```
for (iEdge = 0; iEdge < geometry->GetnEdge(); iEdge++) {
{
    //Gather data from end-points of iEdge
    iPoint = geometry->edge[iEdge]->GetNode(0);
    jPoint = geometry->edge[iEdge]->GetNode(1);

    ProjVel_i = ProjVel_all[iPoint];
    ProjVel_j = ProjVel_all[jPoint];
....
```



**ONERAM6** mesh

#### Big jumps in jPoint

iEdge = 0, iPoint = 0, jPoint = 3731 iEdge = 1, iPoint = 0, jPoint = 83 iEdge = 2, iPoint = 0, jPoint = 1 iEdge = 3, iPoint = 0, jPoint = 84



# Computational Intensity of SU2 FV and DG-FEM Solvers

Computational / Arithmetic Intensity:

- <u>SU2 FV solver (Roe+limiter+gradient</u> reconstruction) ~ 1.25 flops / memory reference
- <u>SU2 DG-FEM solver (p=4, hexahedra)</u> ~ 14 flops / memory reference

## Remember This Slide?



- Only way in modern hardware to get a significant % of peak is through algorithms with higher arithmetic intensity
- Difference can be 5-10% to 50-90% of peak

## ALCF Aurora Early Science Program

Collaboration Stanford (Lele, Alonso) – Argonne Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling





Hydrodynamic phenomena in ICF capsules, showing flow instabilities

Aircraft in transonic buffet conditions. Source: Brunet & Deck, 2008.

Detailed flow physics must be resolved: Ideally suited for higher-order schemes Stanford University

## Pros and Cons of the Methods

	Complex geometries	$\begin{array}{l} \mbox{High-order accuracy} \\ \mbox{ and } hp\mbox{-adaptivity} \end{array}$	Explicit semi- discrete form	Conservation laws	Elliptic problems
FDM	×	~	~	~	~
FVM	~	×	~	~	(√)
FEM	1	~	×	(√)	1
DG-FEM	1	$\checkmark$	$\checkmark$	1	(√)

- $\times$  : Not suited
- ✓ : Suited
- (✓) : Suited, possibly with modifications, but not the most natural (or efficient) choice

Source: Hesthaven and Warburton, 2008

### 2<sup>nd</sup> order: FVM is the best choice High order: DG-FEM is the best choice

### Nodal DG-FEM: Basic Principles (1)

Element 
$$k: U(x_i) = \sum_{j=1}^{N_p} U_j^k \varphi_j^k(x_i)$$



Hyperbolic system of PDE's Weak formulation



$$\frac{\partial U}{\partial t} + \frac{\partial F_i}{\partial x_i} = 0 \quad \Rightarrow$$

$$\iint_{V_k} \frac{\partial U}{\partial t} \varphi_m^k \, dV - \iint_{V_k} F_i \frac{\partial \varphi_m^k}{\partial x_i} \, dV + \oint_{\partial V_k} F_i n_i \varphi_m^k \, d\Omega_k = 0, \quad m = 1, \dots, N_p$$

Nodal DG-FEM: Basic Principles (2)

Contribution from the contour integral

Solution at the interfaces is multiply defined and discontinuous

$$\oint_{\partial V_k} F_i n_i \varphi_m^k \, d\Omega_k \Longrightarrow \oint_{\partial V_k} F_i (U_L, U_R) n_i \varphi_m^k \, d\Omega_k$$

<u>Riemann problem</u>: Any approximate Riemann solver can be used => stabilizes the discretization

1<sup>st</sup> order DG-FEM equals 1<sup>st</sup> order FVM!!!

Nodal DG-FEM: Basic Principles (3)

Nonlinear equations and/or curved elements

### Integrals must be computed with high-order quadrature rules to avoid aliasing. Expensive!!!

Diffusion problems, 2<sup>nd</sup> derivatives

Discontinuous basis functions not suited => must be repaired Even more expensive!!!

However: most operations are local to an element. Extremely well-suited for modern computer hardware

## Implementation in SU2

- Framework of SU2 is very flexible => high-level data structures can be used for any solver, so also DG-FEM
- Input parameter structure can be reused entirely
- FVM parallel I/O functionality could be reused (after some modifications)
- Still a lot of work was required for other low level functions
  - Partitioning of the grid (element wise)
  - Preprocessing is completely different from FVM
  - Standard elements and standard orientation of elements are introduced
  - Spatial discretization is completely new
- But can also reuse entire discrete adjoint formulation to immediately obtain sensitivities of time-accurate (and multi-physics) Qols

## Performance Optimization (1)

- Work still ongoing but nearly complete
- Target architectures: Intel Knights Landing (MIC architecture)
   Intel Xeon
- Hybrid MPI/OpenMP parallelization
- MPI
  - Domain decomposition (one complete halo layer of elements)
  - Overlap computation and communication (including ADER scheme)
  - Use of persistent communication
- OpenMP
  - Aim: Parallelization at for-loop level (sufficient for tests on Xeon)
  - Current data structures are designed for this approach
- Optimized BLAS/LAPACK/LIBXSMM functions must be used to get good performance for matrix multiplications
  - Large contiguous chunks of memory for data storage (not the case for the FVM solver)

### Performance Optimization (2)

Motivation for hybrid parallelization approach Flat MPI does not seem to work too well on KNL Early results on Theta (Argonne)



Elem/Core	122	244	407	1.22K	4.88K	19.5K	39.1K
DOF/Core	4.27K	8.55K	14.2K	42.7K	0.17M	0.68M	1.37M
Efficiency(%)	58.4	70.6	86.6	91.0	99.2	100	N/A Si

### Performance Optimization (2)

Motivation for hybrid parallelization approach Flat MPI does not seem to work too well on KNL Early results on Theta (Argonne)



Number of Cores

Elem/Core	4	8	15	30	960
DOF/Core	31	62	125	250	8000
Efficiency(%)	46.3	63.9	80.6	83.1	N/A

## ADER-DG Schemes Add Complexity...

- MPI Load balancing attempts to make sure every partition has "similar" amount of work
- But exact load balance cannot be achieved a priori
- Solution is to create a list of tasks
  - Volume integrals
  - Surface integrals
  - Multiple steps
  - Varying numbers of elements/faces
- And assign them to threads as work is completed.
- More advanced techniques are also possible



## Performance Optimization (3)

Empirical estimation of workload across different element types for Navier-Stokes simulation

Element Type	Polynomial Order					
	p1	p2	p3	p4		
Tetrahedron	1.000	3.003	6.036	10.071		
Hexahedron	2.291	6.655	14.621	21.033		
Prism	1.620	8.029	17.169	37.098		
Pyramid	0.811	5.847	14.609	13.195		

- Accurate estimate of workload for each element type is necessary for load balancing, especially in ADER implementation
- Time accurate local time stepping (ADER) uses load balancing that is dependent on the speed of execution of volume and surface computations in each type (and order) of elements



# **OpenMP Tasks in Regent**

- Legion now supports OpenMP tasks
  - Realm implements OpenMP ABI compatible to ICC/LLVM and GCC



## Performance Optimization (4)

Speed-up by using intel MKL



- In initial testing, use of libXSMM provides additional performance improvements
- Tests ongoing on Theta

### Strong Scaling Test – 3D SD7003 Airfoil

Strong scaling test - SD7003, 0.49M Elem, p = 4, Hexs



### **DG-FEM Solver Performance on KNL**





## What About Managing the Numerical Error?

- Functional-based error estimation / adaptation theory has been around for a while
- In a few years this problem will be solved and commercial solvers will incorporate these capabilities





- We have teamed up with INRIA and are using their mesh adaptation library
- Will we also be able to manage the uncertainties simultaneously? This is the grand challenge





## Conclusions & Ongoing Work

- Optimized CFD solvers for new computing platforms require significant effort and expertise
- Scientists and engineers should focus on doing engineering and science: turbulence modeling and others
- Performance / scalability optimization expertise is quickly disappearing in academe
- Multi-disciplinary teams are needed...and hard to come by
- Must amortize the effort over a larger base of users
- Community codes may be an answer: high-performance out of the box







# **TUDelft** UNIVERSITY OF TWENTE.



## **Stanford** | ENGINEERING

Aeronautics & Astronautics





















The Open-Source CFD Code