Co-Design of Hardware/Software for Predicting MAV Aerodynamics





E. de Sturler, **W. Feng**, C. Roy, A. Sandu, D. Tafti


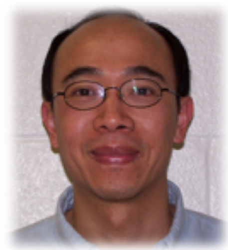


J. Edwards, H. Luo, F. Mueller

Fariba Fahroo, Computational Mathematics

Vision

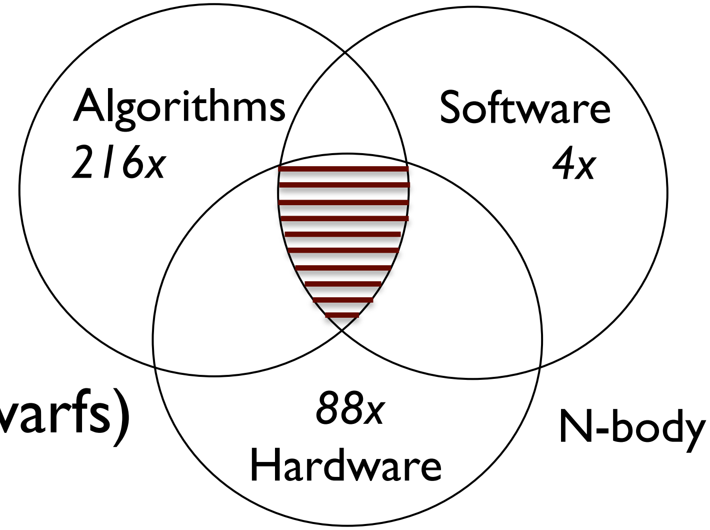

- Synergistic co-design process for the *structured/unstructured grid motifs* (or dwarfs) in computational fluid dynamics (CFD) to support aerodynamic predictions for micro-air vehicles (MAVs).
  - Malleable **algorithms**
    … that can be mapped and optimized in **software**
    … onto the right type of processing core in **hardware**
    … at the right time
  - Co-design feedback to vendors to assist in guiding future hardware design

Synergistic Co-Design: Enabling and Empowering

Team @ VT & NCS

CS : Computer Science
Math : Mathematics
AOE/MAE : Aerospace & Ocean Engg/Mechanical & Aerospace Engg.

- Virginia Tech (5 PIs + 7 students/staff)
  - Eric de Sturler — Numerical Methods (solvers & preconditioners)
  - Wu Feng, Lead PI — Parallel Computing (performance, programmability, portability)
  - Chris Roy — CFD (structured grid and ALE mesh movement)
  - Adrian Sandu — Numerical Methods (time stepping & discretization)
  - Danesh Tafti — CFD (pressure-based multiblock structured)
  - Research Scientist (1), Postdocs (1), and Graduate Students (5)
- North Carolina State University (3 PIs + 3 students/staff)
  - Jack Edwards — CFD (multiblock structured w/ implicit solvers)
  - Hong Luo — CFD (unstructured grid / compressible)
  - Frank Mueller — Parallel Computing (languages, compilers, scalability)
  - Postdocs (1), Graduate Students (2)

Interdisciplinary Collaboration

(via formal biweekly meetings + informal meetings)

CS : Computer Science
Math : Mathematics
AOE/MAE : Aerospace & Ocean Engg/Mechanical & Aerospace Engg.

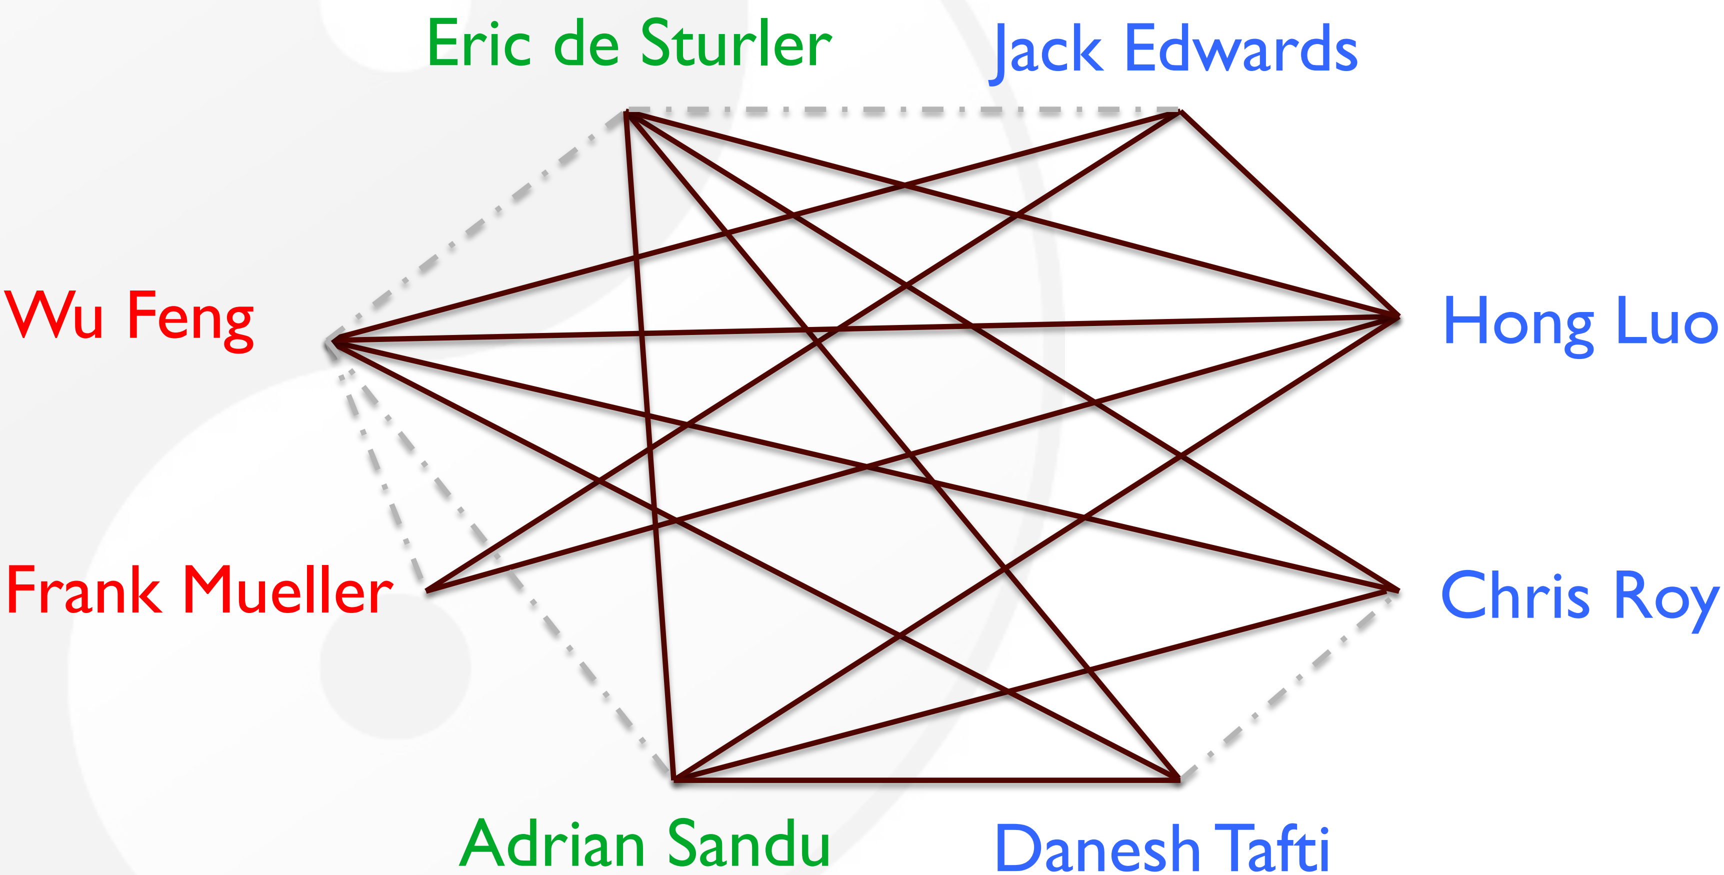

See appendix for details of collaborations

Why Synergistic Co-Design? Why Now?

- Increasing heterogeneity in computing resources

Altera FPGA

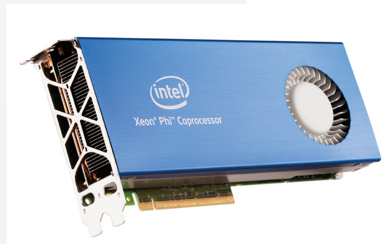

Intel MIC

NVIDIA GPU

TI DSP

… across a wide variety of environments

VirginiaTech
Invent the Future

W. Feng, wfeng@vt.edu, 540.315.1545

Heterogeneous Systems in HPC

- Statistics
  - *Four* out of top 10 systems
  - Performance share in Top500 systems
    5% (2009) → 35% (2013)
- HokieSpeed
  - CPU+GPU heterogeneous supercomputer with large-scale visualization wall
  - Debuted as the *GREENEST* commodity supercomputer in the U.S. in Nov. 2011

Tianhe-2

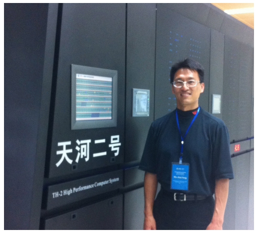

HokieSpeed Viz Wall
(Eight 46" 3D HDTVs)

THE GREEN 500

Top 10 Source: top500.org

1. **Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 [illegible]ess-2, Intel Xeon Phi 31S1P
   NUDT
2. **Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x
   Cray Inc.
3. **Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom
   IBM
4. K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect
   Fujitsu
5. **Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom
   IBM
6. **Piz Daint** - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x
   Cray Inc.
7. **Stampede** - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P
   Dell
8. **JUQUEEN** - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect
   IBM
9. **Vulcan** - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect
   IBM
10. **SuperMUC** - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR

HOKIESPEED

A GPU-Accelerated Supercomputer for the Masses

- Hardware
  - Total Nodes: 209, where each compute node consists of
    - Motherboard: Supermicro 2026GT-TRF Dual Intel Xeon
    - CPUs: Two 2.4-GHz Intel Xeon E5645 6-core (12 CPU cores per node)
    - GPUs: Two NVIDIA Tesla Fermi GPUs (M2050/C2050)
- Software
  - CUDA 4.1.28, 5.0.35, 6.5.14 via module load/unload
  - PGI compilers v13.4 w/ CUDA Fortran & OpenACC support
  - OpenMPI & MVAPICH2-1.9
  - Torque/Maui job management system

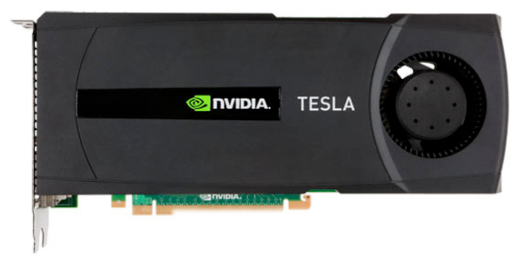

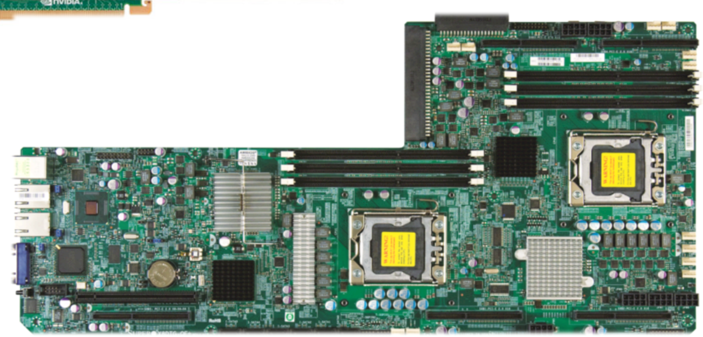

ARC Cluster

- Hardware
  - 2x AMD Opteron 6128 (8 cores each)
    - 108 nodes = 1728 CPU cores
  - NVIDIA GTX480, GTX680 , C2050, K20c: 108 GPUs
  - Mellanox QDR InfiniBand: 40Gbit/s
- Software
  - CUDA 5.0
  - PGI Compilers V13.4 w/ CUDA Fortran & OpenACC support
  - OpenMPI & MVAPICH2-1.9 w/ GPUDirect V2 capability
  - Torque/Maui job management system

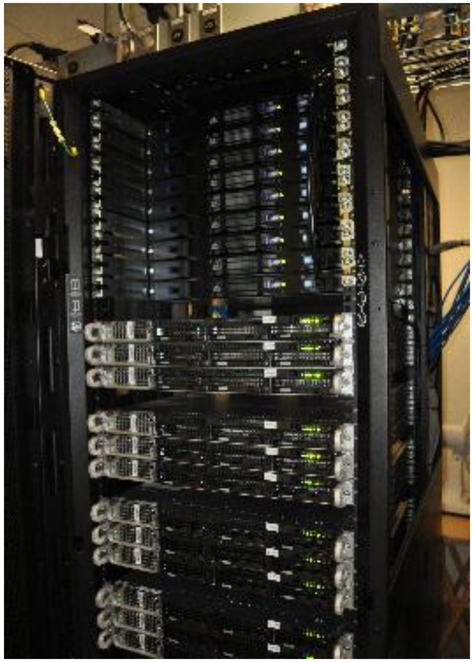

Diversity of Heterogeneous Systems

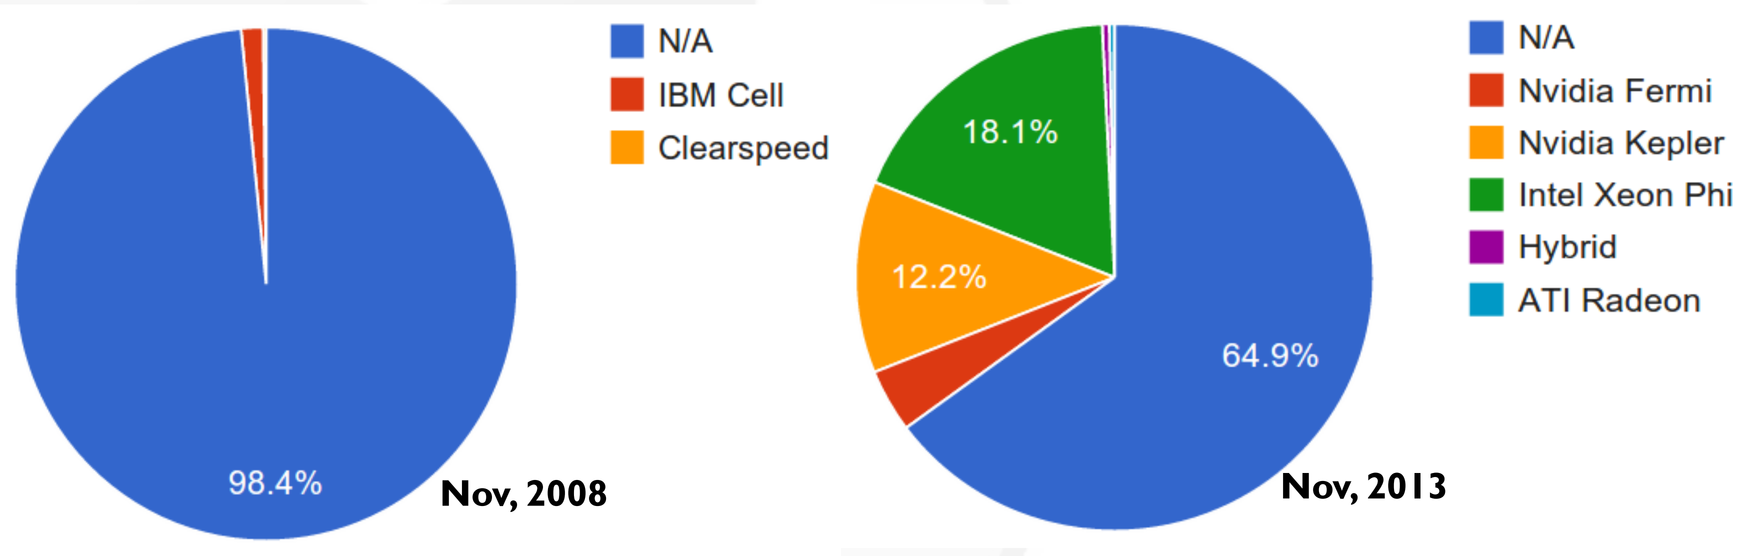

Performance Share of Accelerators in Top500 Systems

Layered Co-Design

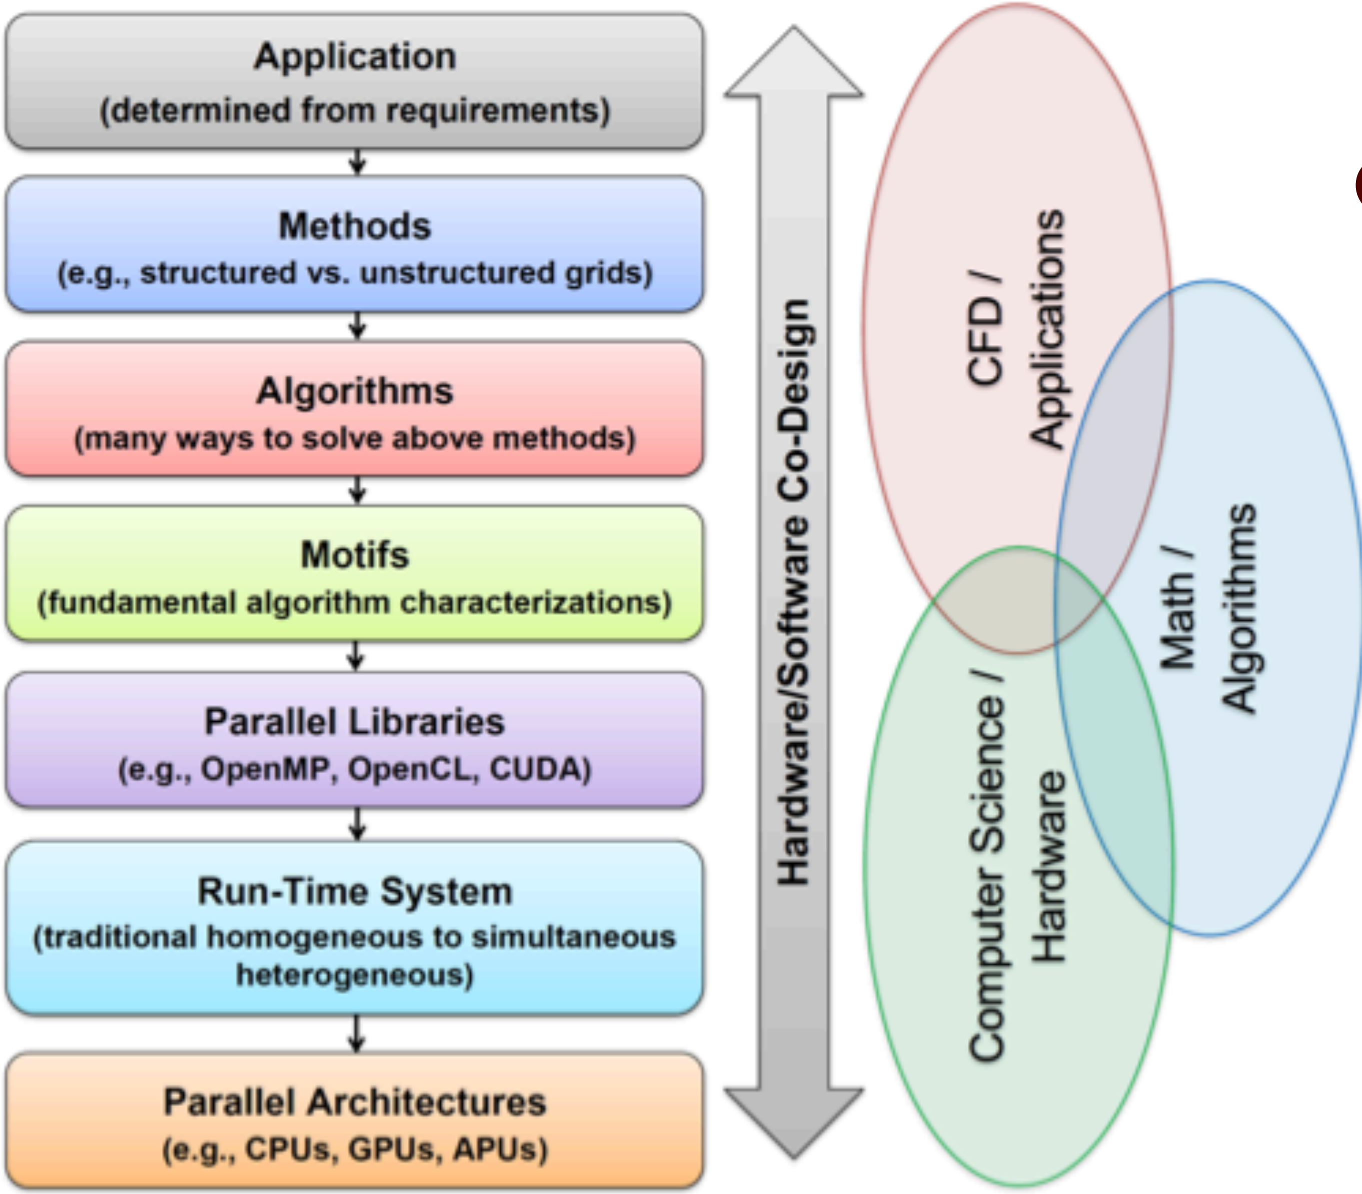

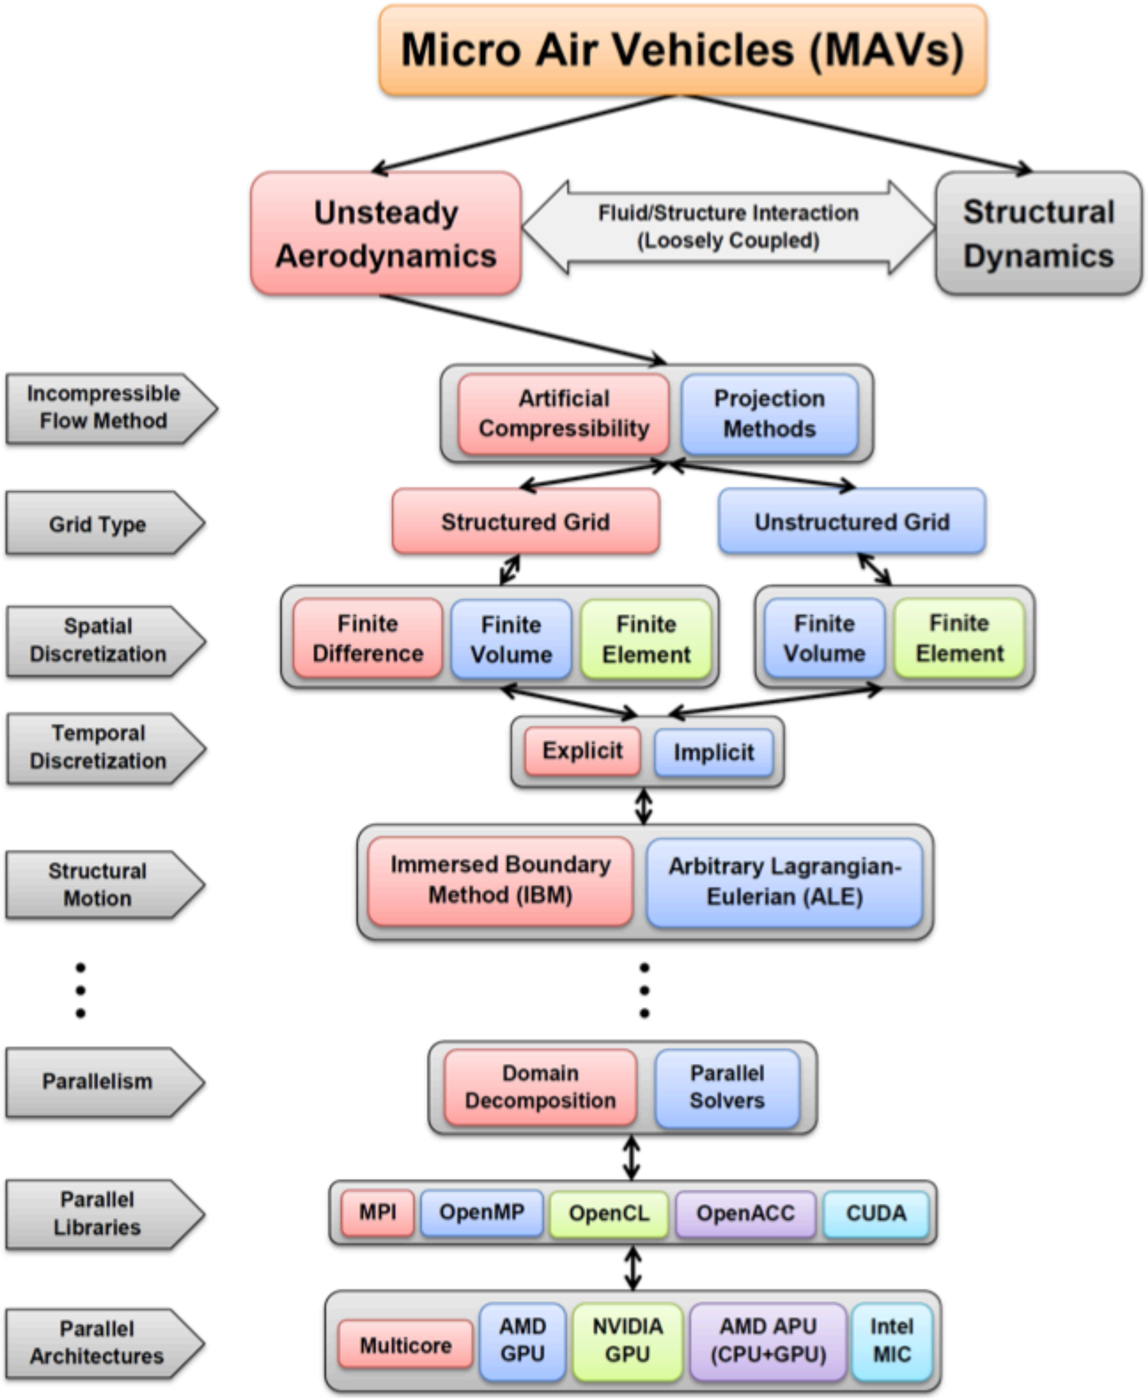

Co-Design for Micro Air Vehicles (MAVs)

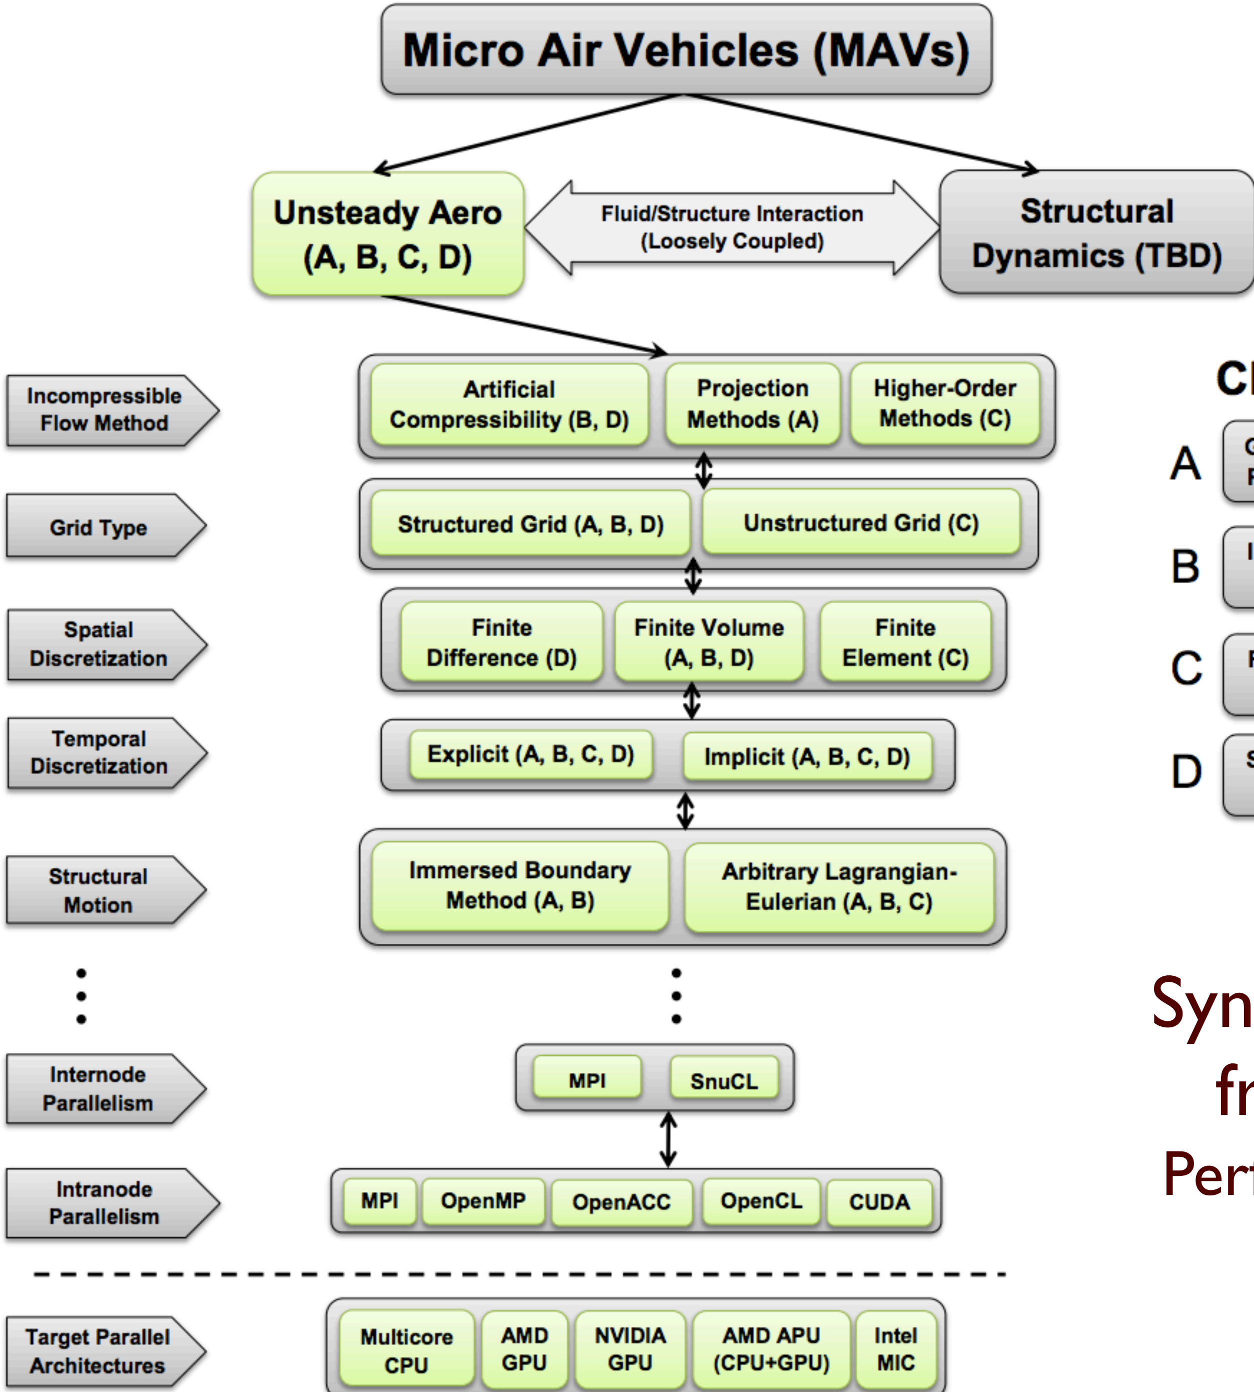

CFD Codes

- **A** — GenIDLEST: Struct, Press Project, FVM
- **B** — INCOMP3D: Struct, Art Compr, FVM
- **C** — RDGFLO3D: Unstr, FEM
- **D** — SENSEI: Struct, Art Comp, FVM

Synergistic Co-Design from 10,000 Feet:
Performance Perspective

SyNeRG
synergy.cs.vt.edu

Bat Wing Simulation for GenIDLEST Code

Amit Amritkar and Danesh Tafti, Dept. of Mechanical Engineering, Virginia Tech

GPU

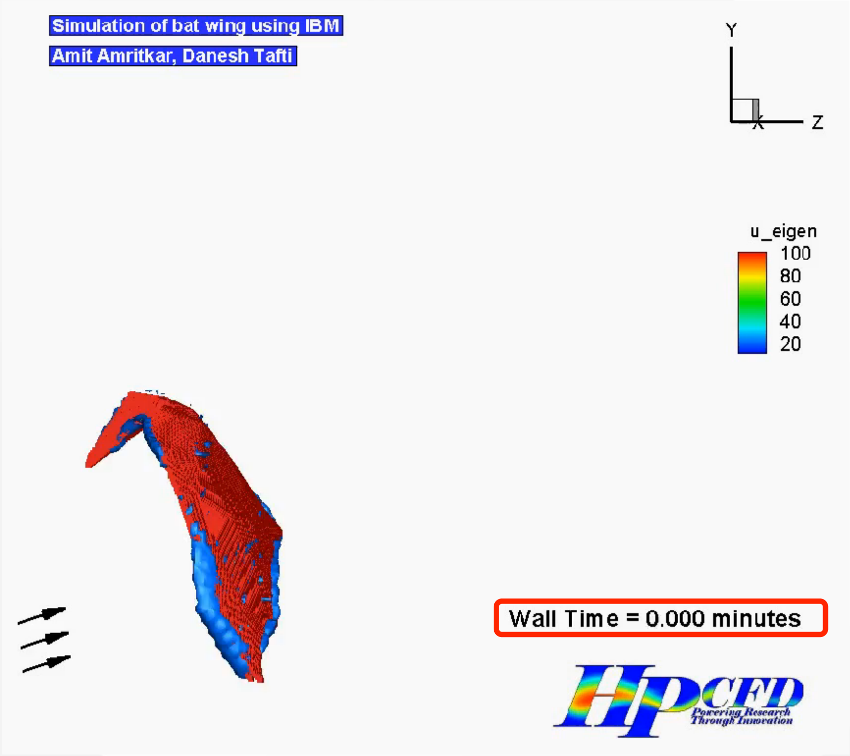

CPU

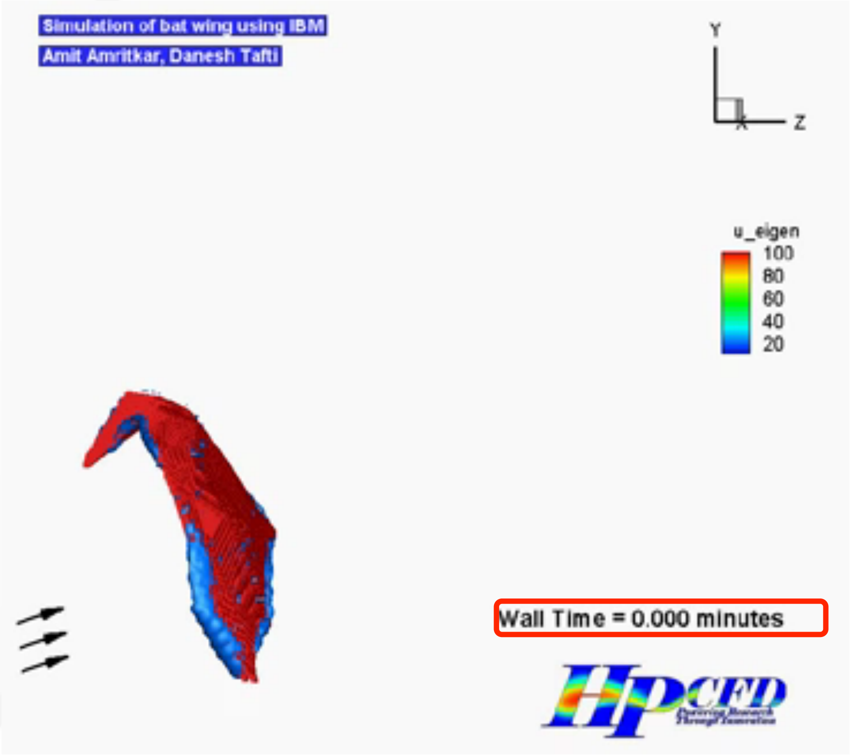

Co-Design Around Three P's: Performance, Programmability, Portability

"Productivity = Performance + Programmability + Portability"

– Multi-dimensional optimization across two or more P's

… first *manual co-design* … then *automated co-design*

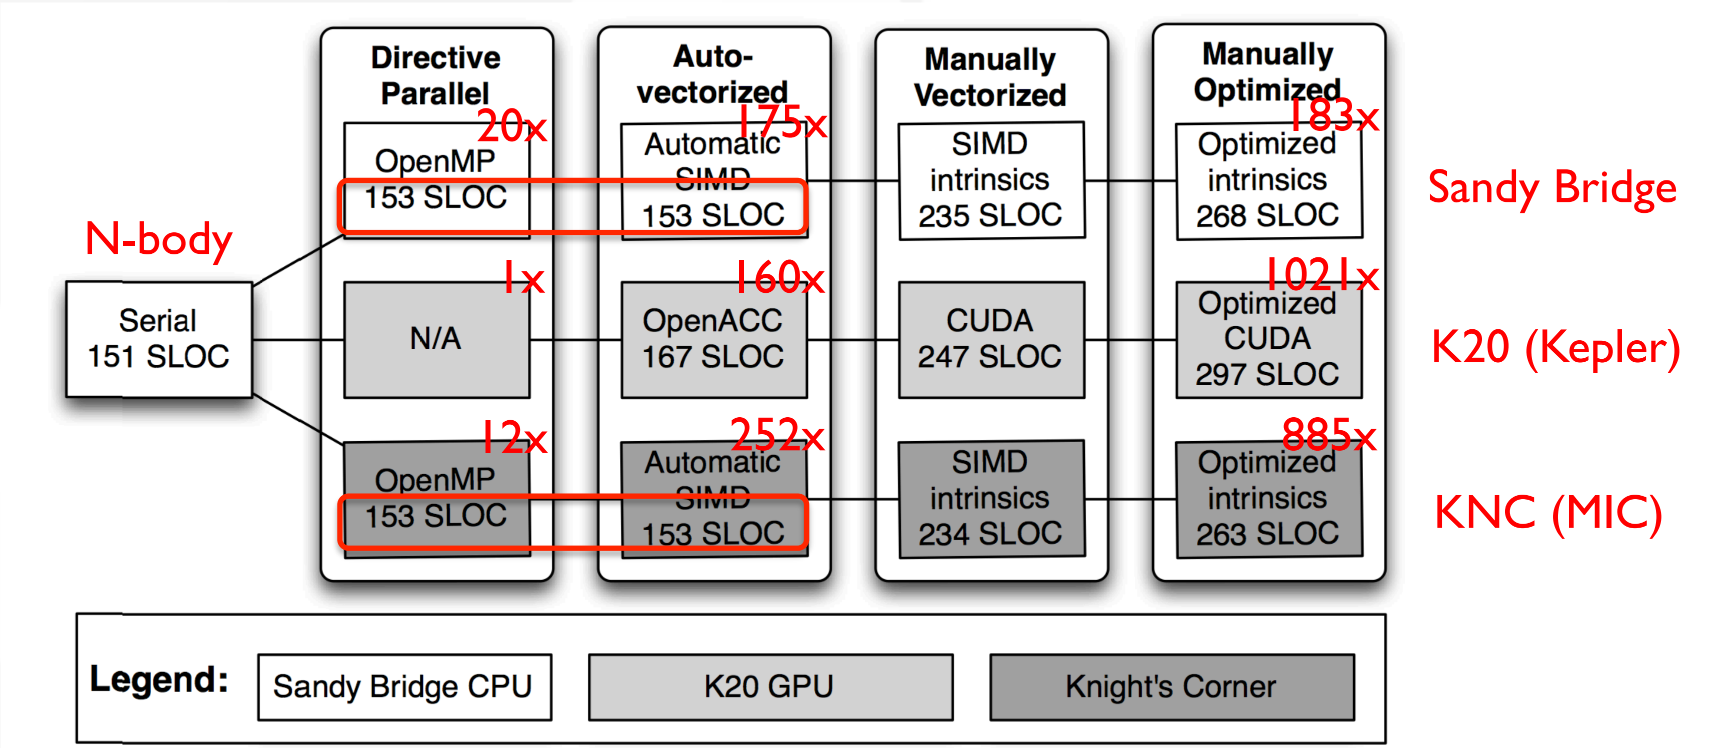

What's Next?

- Two-Year Renewal Option for Years 4 and 5
- Brainstorming …
  - Continue with the status quo?
    - Manual Co-Design for Each Code + Starchart + Stargazer
    - Multi-Dimensional Performance, Programmability, and Portability
      - Platforms, languages, run-time systems
    - Continued Extensions to MetaMorph
    - MetaMorph-Enabled CFD Codes
  - Explore (or add) new directions
    - Multi-Pronged Programming Language w/ Auto-Optimization Approach
    - Adaptive Run-Time System for Heterogeneous Computing
    - MetaMorph + ParMETIS? How? (Amit A.)
  - Something more radical?

Acknowledgements

- This work was funded by the Air Force Office of Scientific Research (AFOSR) Computational Mathematics Program
  - Program Manager: Fariba Fahroo
  - Grant No. FA9550-12-1-0442