Task 4:Machine Learning to Accelerate CFD Models

Dirk Van Essendelft, Terry Jordan, Mino Woo, Tarak Nandi

Project Objective

Research Goal

Build an advanced collaborative framework specifically targeted towards CFD on the most advanced HPC/AI hardware with native support for AI and ML algorithms

Aligned with FE Objectives

Increasing computational speed without sacrificing accuracy will directly supports:

- Modernization of existing coal plants
- development of coal plants of the future
- Reduction of the cost of carbon capture, utilization, and storage (CCUS)

Project Origins

Driving Question

MFiX + TensorFlow

Can we write MFiX in TensorFlow so that we can create a single, unified framework for doing both CFD and AI/ML on emerging hardware designed for AI/ML?

- TensorFlow is the most used AI/ML framework
- TensorFlow has a simple API and allows for both surface level hardware agnostic coding and the ability to deeply optimize hardware specific implementations if needed
- Get speed boosts from AI/ML hardware
- Get speed boost from AI/ML accelerated algorithms
- Simplify implementation of AI/ML models in MFiX

Current Status

Where are we now?

- In Second Full Year of Development, First year in CARD
- Ahead of schedule with Milestones
  - EY20.4.A A demonstration of a multidevice linear solver relative to the existing single device solver (9/30/2020)
  - EY20.4.B A demonstration of a granular simulations using the TensorFlow based solver. (9/30/2020)
  - EY21.4.C A demonstration of a simple fluid bed simulation in the TensorFlow based solver. (3/30/2020)
- Have a functioning, coupled MP-PIC code implemented in TensorFlow
  - Solves all transport equations on available devices followed by a multi-device solve of continuity
  - Ready to accept AI/ML models
  - Does not yet support complex geometries
  - Does not yet support energy, species, or reactions

Comparison to State of The art

MFIX AI on GPU hardware: MP-PIC Solver

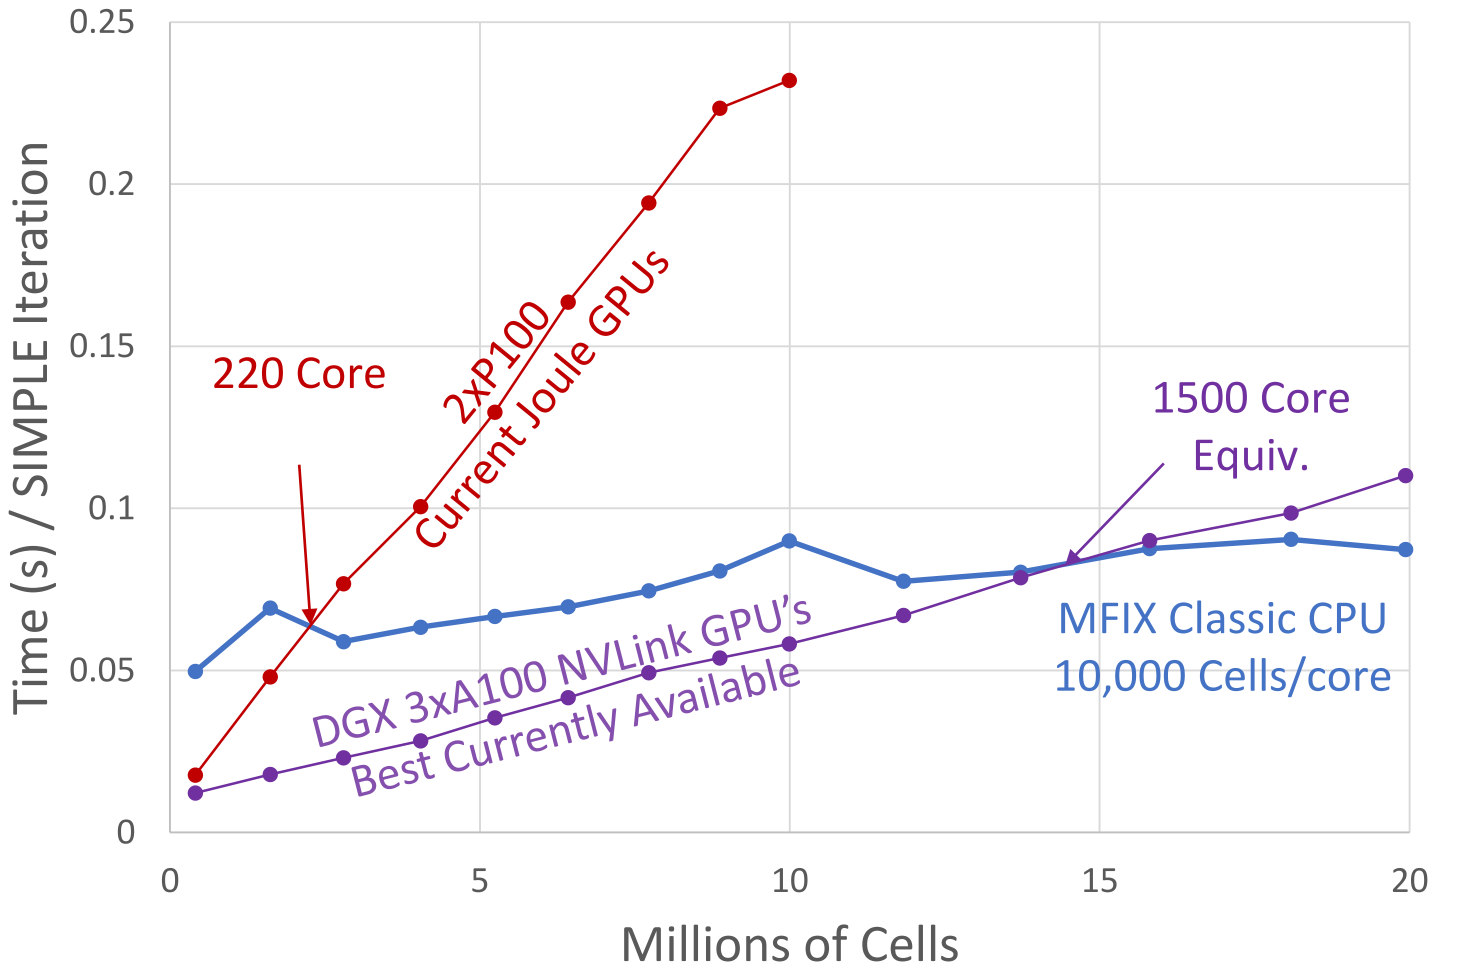

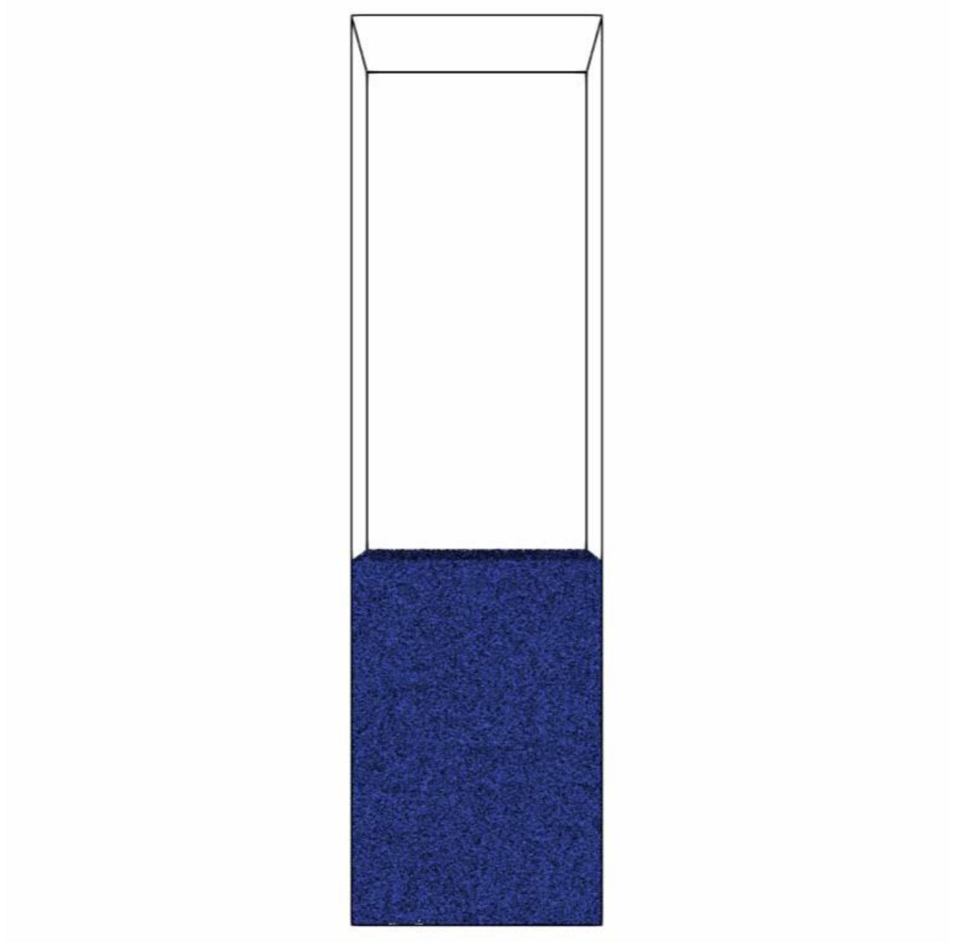

3.8 Million Cells, 16.8 Million Parcels

Comparison to State of The art

MFIX AI on GPU hardware: MP-PIC Solver

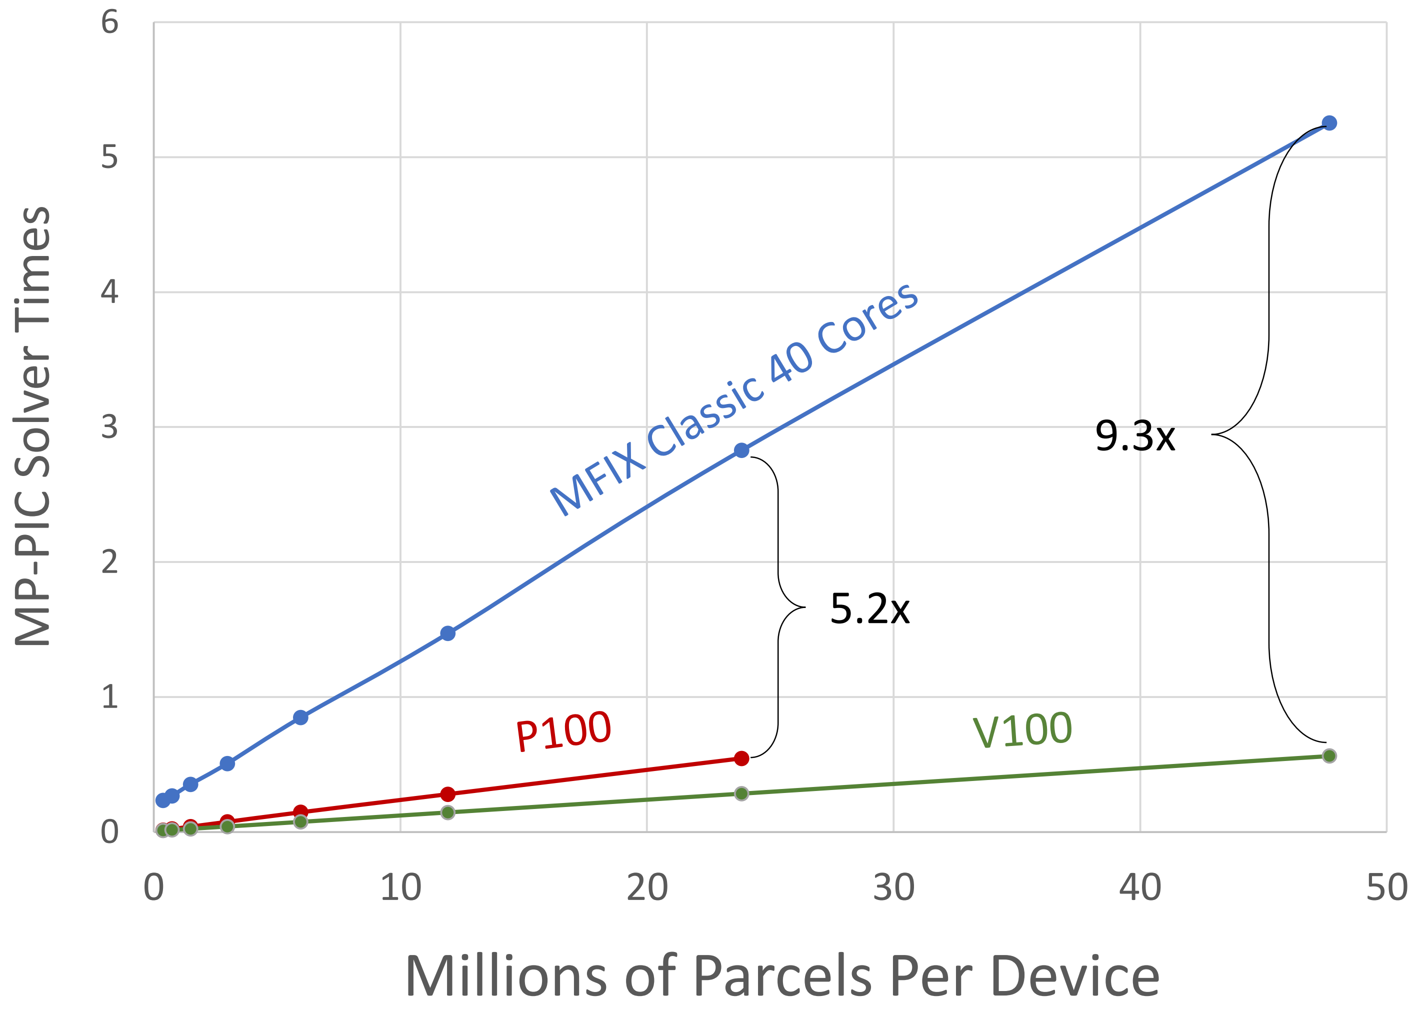

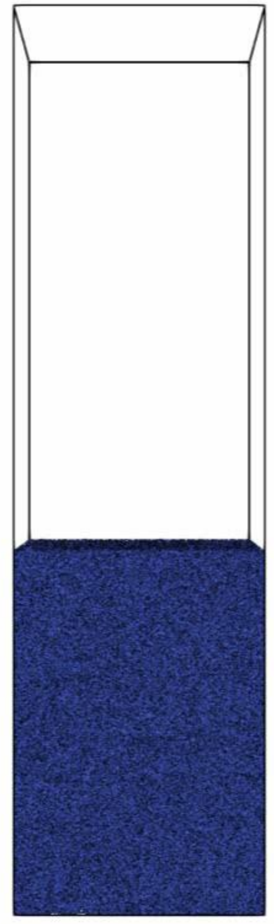

3.8 Million Cells, 16.8 Million Parcels

CFD on Emerging Hardware

MFIX (linear Solver) on Joule vs equivalent linear solver on Cerebras CS1

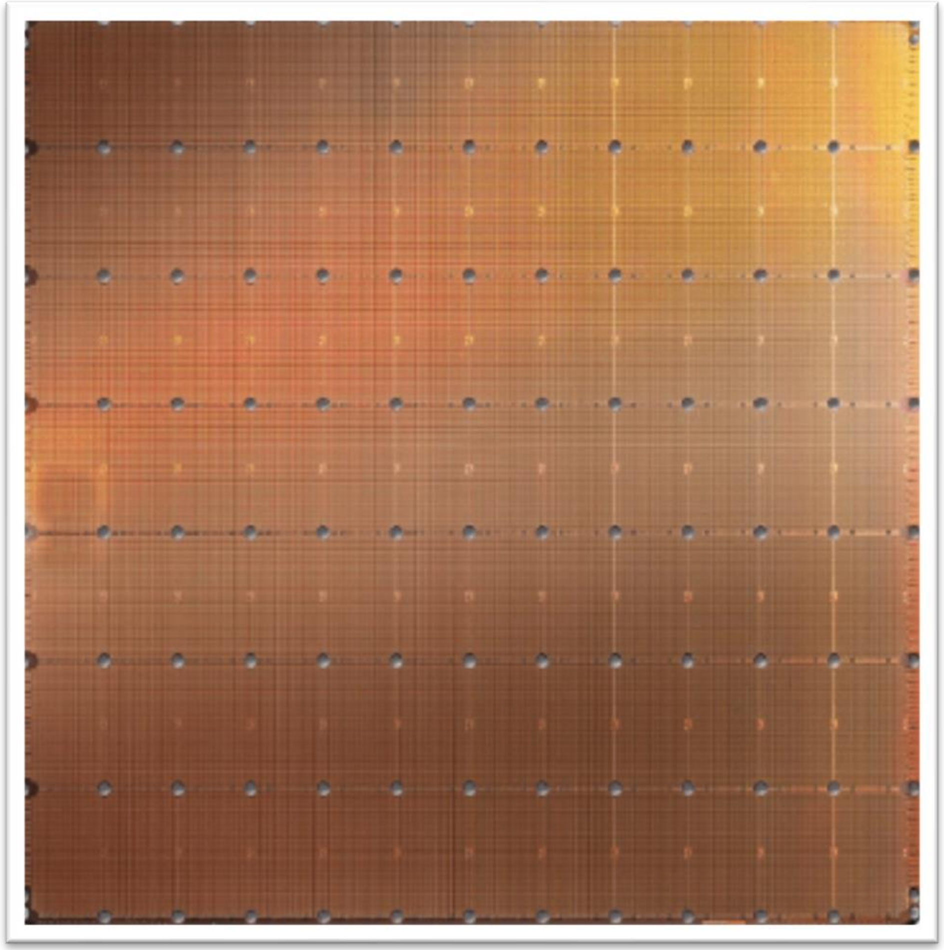

cerebras.net

Size	462 cm^2
Cores	400,000
Transistors	1.2T
Memory Band Width	9.6 PB/s
Interconnect Bandwidth	100 Pb/s
Memory	18GB
Power	20kW

Kamil Rocki, Dirk Van Essendelft, Ilya Sharpov, Robert Schreiber, Michael Morrison, Vladimir Kibardin, Andrey Portnoy, Jean Francois Dietiker, Madhava Syamlal, and Michael James, 'Fast Stencil-Code Computation on a Wafer-Scale Processor', in *SC'20 Proceedings of The International Conference for High Performance Computing, Networking, Storage, and Analysis* (Atlanta, Ga: IEEE Press, 2020).

Cerebras CS1: Unrivaled Speed

MFIX (linear Solver) on Joule vs equivalent linear solver on Cerebras CS1

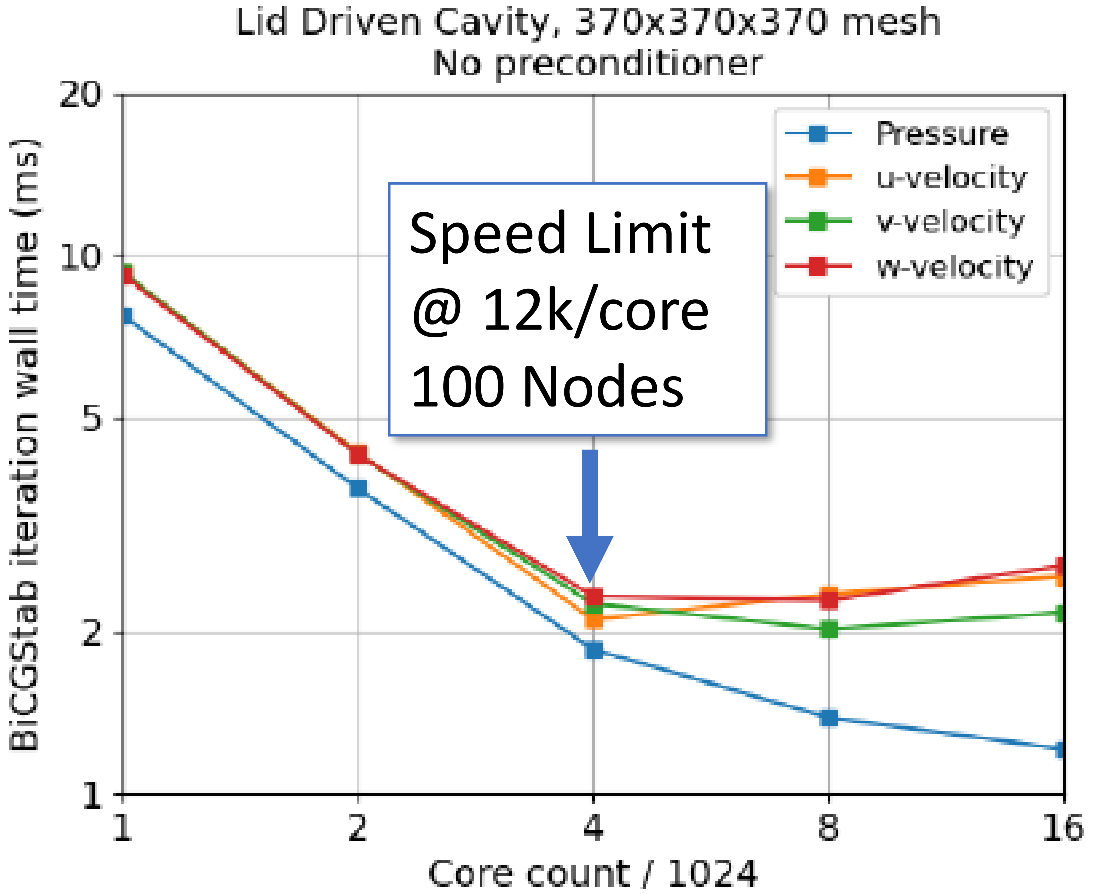

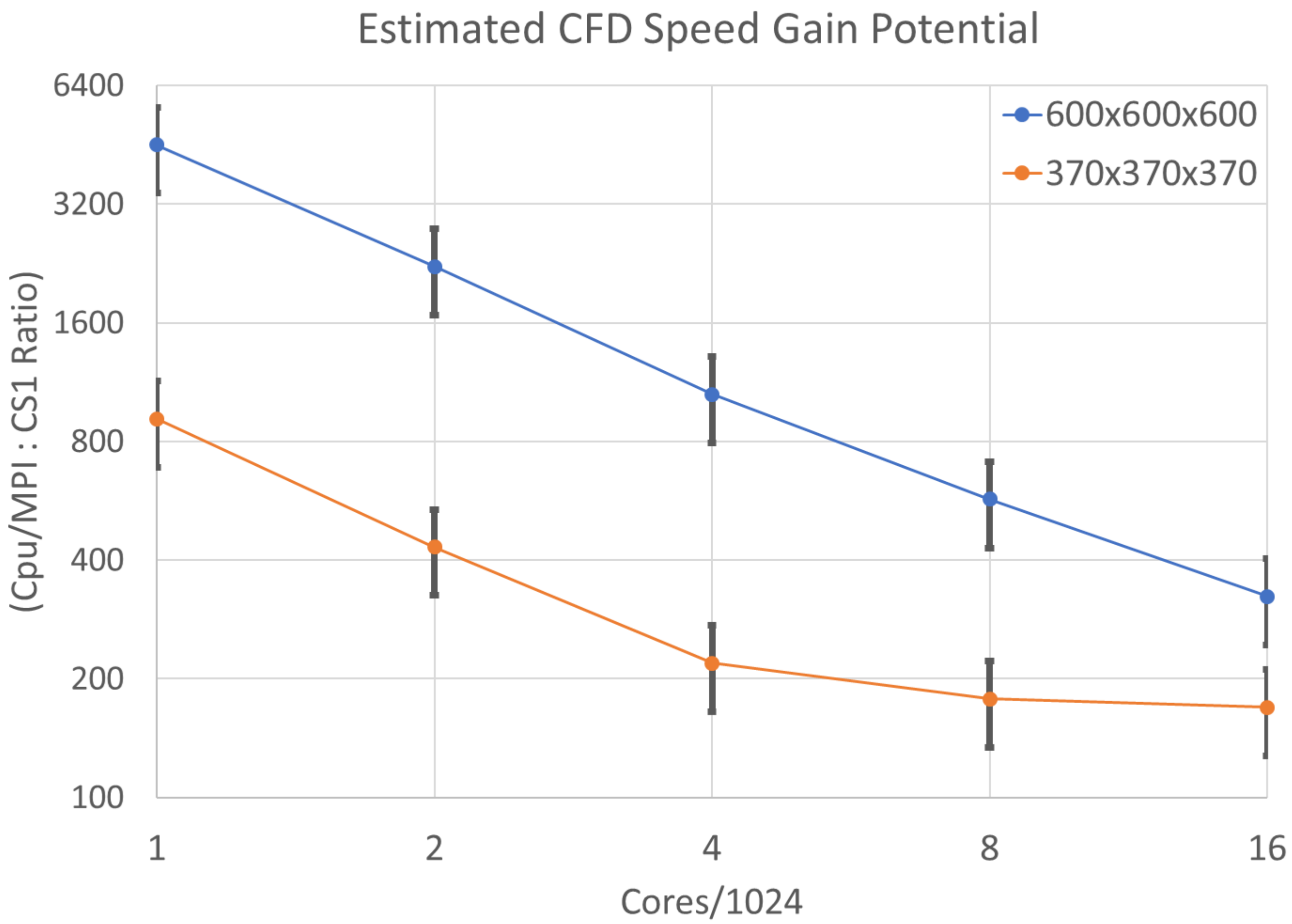

Kamil Rocki, Dirk Van Essendelft, Ilya Sharpov, Robert Schreiber, Michael Morrison, Vladimir Kibardin, Andrey Portnoy, Jean Francois Dietiker, Madhava Syamlal, and Michael James, 'Fast Stencil-Code Computation on a Wafer-Scale Processor', in *SC'20 Proceedings of The International Conference for High Performance Computing, Networking, Storage, and Analysis* (Atlanta, Ga: IEEE Press, 2020).

Preparing Project for Next Steps

Market Benefits/Assessment

- While CFD is very powerful design tool, the extremely high computational cost limits its practical application
  - 9x improvement on PIC, up to 4x on fluid with current gen hardware with TensorFlow
  - >200x on emerging hardware
  - Makes CFD for optimization and UQ tractable
  - Could open up new application areas in real time or faster than real time CFD

Technology to Market Plan

- Make incremental releases on the existing MFiX platform so existing MFIX users can pick up and run the tools for their FE supported work
- Remaining needs: complex geometry, species/energy, reactions
- Potential New Research: DMP method within TensorFlow, Cerebras CS1
- Existing/Potential Collaborators:
  - Industrial relationships with NVIDIA and Cerebras
  - Chris Guenther with TTNEP FWP leveraging frame work with WVU collaboration
  - FOA 2193 - ML models for non-spherical drag

Concluding Remarks

Important Concepts and Next Steps

- The most important thing that this project brings is very high levels of computational speed and efficiency without sacrificing accuracy.
  - Means more work gets done in less time and at lower costs
  - Directly translates to reduced uncertainty, design time, cost, and risk for FE applications
- Project next steps:
  - Explore DMP method in TensorFlow – could significantly improve scaling
  - Deploy the cut cell methodology for complex geometries
  - Formalize a development plan to take advantage of the unique capabilities on the CS1.
  - Species/Energy, reactions (out years)