SAND2019-7318PE

Exceptional service in the national interest

An Overview of Machine Learning for Scientific and High Performance Computing at Sandia

Justin Newcomer, Manager, Machine Intelligence
Sandia National Laboratories
jtnewco@sandia.gov

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

A Digital Engineering Strategy requires a Data-Centric Foundation

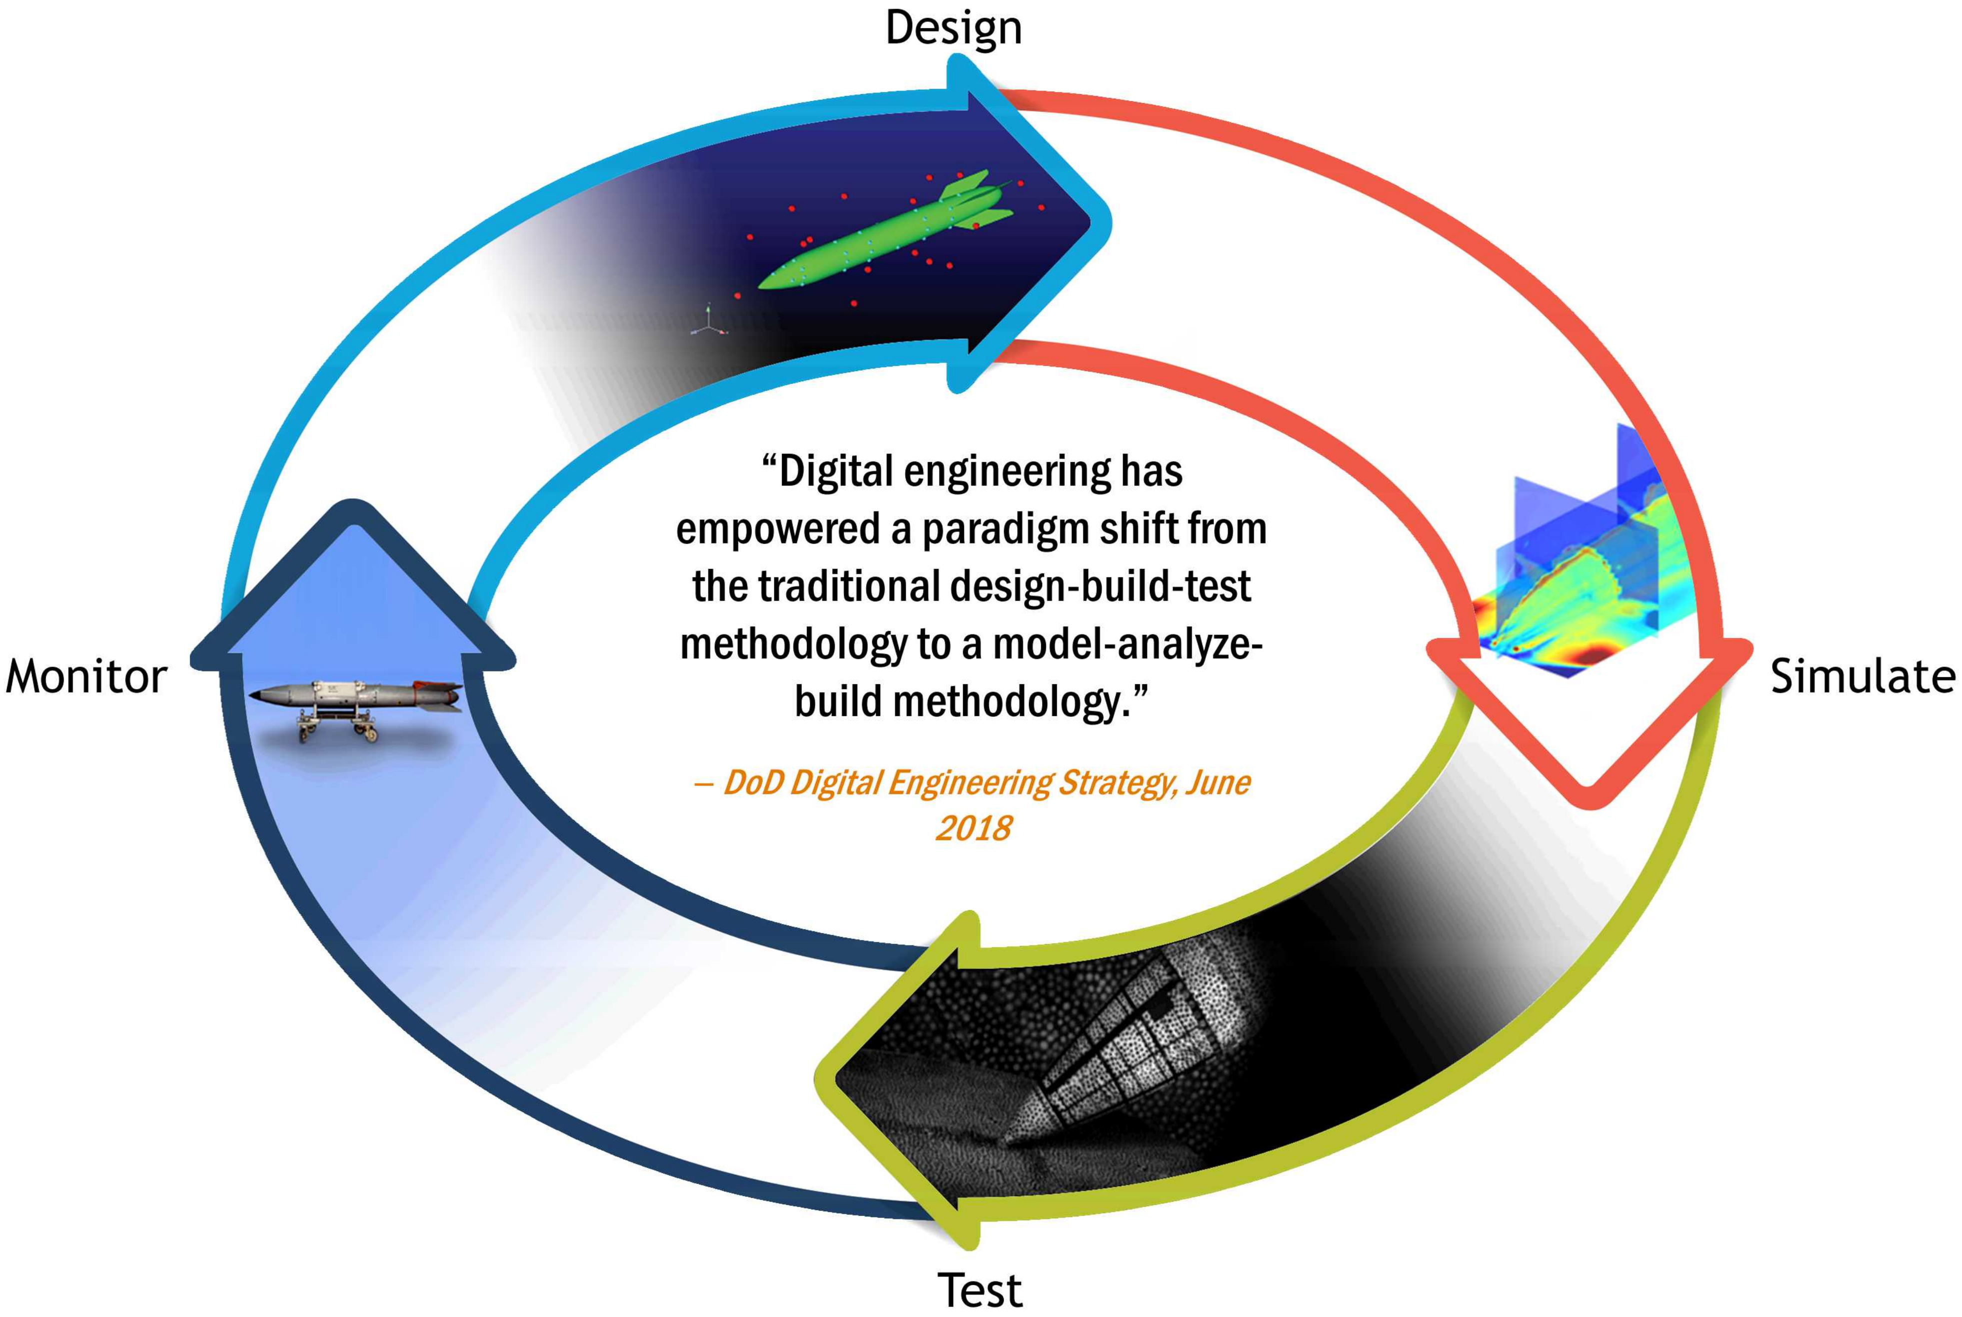

"Digital engineering has empowered a paradigm shift from the traditional design-build-test methodology to a model-analyze-build methodology."

– DoD Digital Engineering Strategy, June 2018

A Data-Centric Foundation will Enable Transformative Goals

Reduce the design cycle time	Increased production throughput	Re-think the surveillance program for the 21st century
DATA STREAM	DATA STREAM	DATA STREAM
Agile, model-based design workflow Rapid cycles of learning	Credible mod/sim results Maintain International HPC/HPDA leadership	Risk-, physics-, and age-aware surveillance by design Synthesis of experimental and mod/sim results for critical decision making

National ASC Advanced Machine Learning (AML) Focus Areas		
Improved efficiency in design process	Data-driven physics models	Enhanced experimental design
Anticipatory stockpile decision making	Reduced computational cost	

National ASC AML Priorities		
Physics-constrained ML	Employ ML with sparse data	Invest in credible ML
Learning HPC hardware systems	Improve data specifications	Build talent at our labs

ML Enables Automatic Mesh Generation

PI: Steve Owens

- **Problem**
  - Geometry preparation and meshing for computational simulation is bottleneck (consuming 70%+ of analyst time)
  - Analyst/engineer must have extensive domain-specific expertise to manage many individual complex problems and tasks
  - Must produce verifiably accurate physics appropriate mesh ready for simulation
- **Technical Approach**
  - Identify tasks currently done by analysts to train machine learning models
  - Capture and label operations performed by expert using existing software
  - Build a feature library of geometric characteristics commonly encountered in CAD models and identify solutions for effectively modifying CAD for best resulting mesh
  - Explore machine learning models that provide best solutions for CAD features with associated solution labels
- **Results/Accomplishments**
  - Developed ML techniques to rank geometry-modification operations by their likelihood of yielding a meshable model
  - Provides insight on which geometric features are most useful for machine leaning, and would be relatively easy to integrate into the analyst workflow if successful

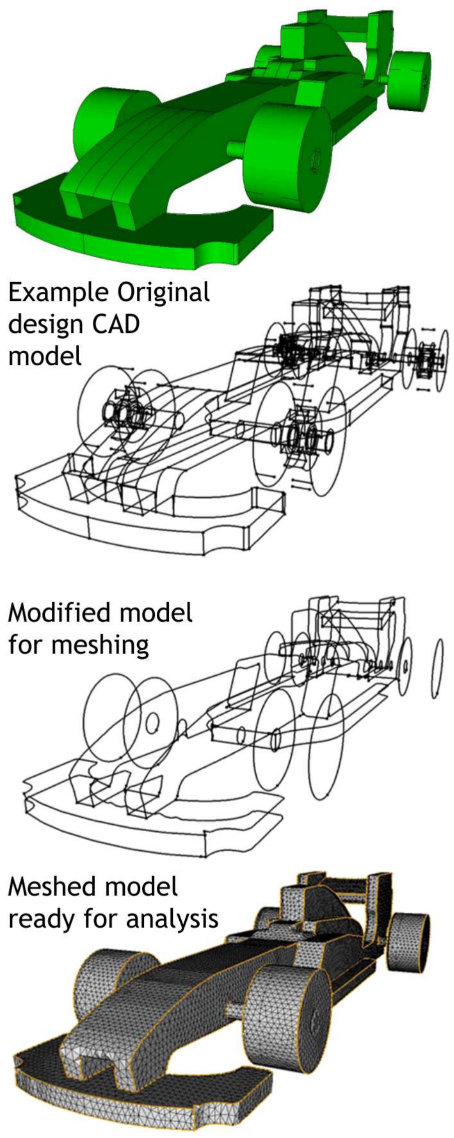

Main Focus Areas: Improved efficiency in design process and Reduced computational cost → Agile model-based design workflow and Rapid cycles of learning → Reduce the design cycle time

Deep Learning Enabled Discovery of Anomalies

PI: Emily Donahue

ML with Sparse Data

- **Problem**
  - It is difficult to train ML networks to identify rare/never-before-seen features with high confidence
- **Technical Approach**
  - Train a neural network (with 3000 images) to understand only "good" features
- **Results/Accomplishments**
  - Deep learning model finds anomalies in seconds that would take a human hours to find
  - Enables automated anomaly detection in production and surveillance

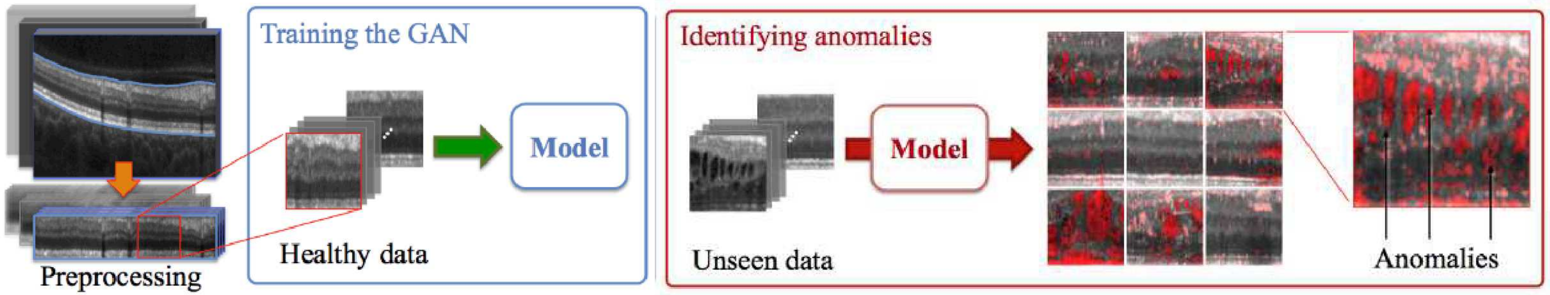

Schlegl et. al. Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker Discovery. 2017

Real query image

Anomaly detection

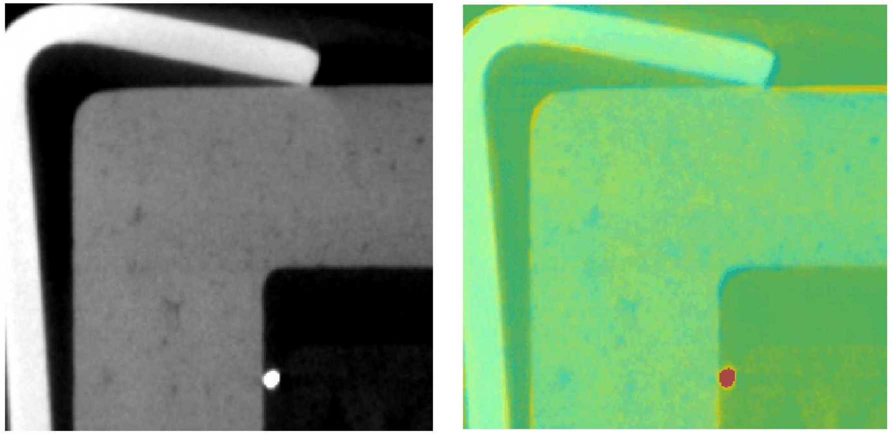

X-Ray CT scan of faulty electronic component

Main Focus Areas: Anticipatory stockpile decision making and Enhanced experimental design

Risk-aware surveillance and Synthesis of experimental and mod/sim results for critical decision making

Increased production throughput

Re-think the surveillance program for the 21st century

ML for Reduced Order Models

PI: Kevin Carlberg

Credible ML

- **Problem**
  - High-fidelity computational physics simulations on HPC systems can take hours or days to execute
  - Lengthy execution time limits the design space explored during conceptual design
  - Need a faster, more efficient means of simulating complex physics problems
- **Technical Approach**
  - Create Reduced Order Model (ROM) from high-fidelity simulation data that
    - Executes faster via dimensionality reduction using autoencoders without significant reduction in accuracy
    - Preserves important physical properties (e.g., conversation laws)
    - Uses Machine Learning Error Models (MLEM) to quantify uncertainty
- **Results/Accomplishments**
  - Reduced order surrogate models and theory have been developed for turbulent flow simulations
  - Runtimes are 100-1000 times faster and are only 1% less accurate than the high-fidelity simulations
  - MLEM can predict errors with validated statistical properties

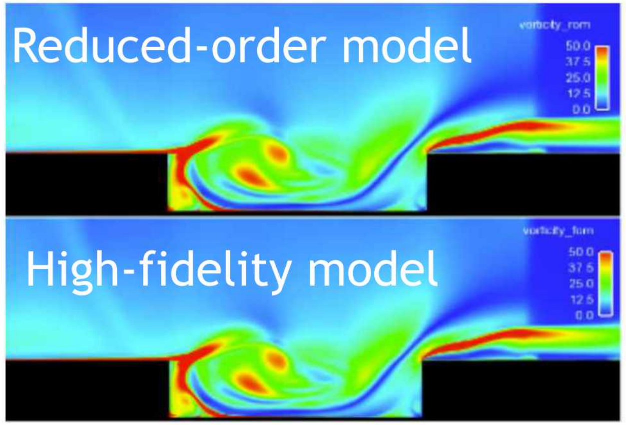

Turbulent flow vorticity field

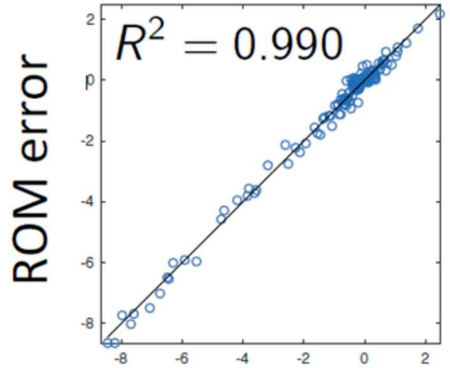

SVM error prediction

Main Focus Areas: Data-driven physics models and Reduced computational cost

Agile model-based design workflow, Credible mod/sim results, and Physics-aware surveillance

Reduce the design cycle time

Increased production throughput

Re-think the surveillance program for the 21st century

Diagnosing HPC Performance Variations

PI: Vitus Leung

Learning Hardware

- **Problem**
  - On HPC systems, the same job on the same system can vary in performance up to 100%
    - Leads to poor scheduling
    - Results in reduced efficiency (costly on large-scale systems)
  - Much of the variations are caused by system anomalies
    - E.g., Shared resource contention, firmware bugs, CPU throttling for thermal control, orphan processes from previous jobs

- **Technical Approach**
  - Created an automated ML framework to detect and classify anomalies
    - Based on resource usage data (CPU, memory, network)
    - Using concise features, low-overhead
    - Generally applicable

- **Results/Accomplishments**
  - Framework outperforms existing methods
    - Evaluated in two different HPC environments, over 0.97 F-Score
  - Easy-to-compute statistical feature extraction
    - Storage overhead reduced to less than 10%
    - Computation overhead below 1% of a single core

Main Focus Areas: Reduced computational cost

Maintain International HPC/HPDA leadership

Reduce the design cycle time

Increased production throughput

Re-think the surveillance program for the 21st century

Learning Hardware

ML for System Software

PI: Craig Vineyard

- **Problem**
  - The growing complexity of computing systems, ranging from cell phones to supercomputers, is becoming difficult for developers to manage
  - More intelligent and automated mechanisms are needed to avoid unintended resource oversubscription and manage the placement and movement of data and computation in extremely heterogeneous systems
- **Technical Approach**
  - We believe that ML can be utilized to infer intelligent approaches to resource management
  - In particular Reinforcement Learning (RL)
    - ML paradigm where an agent makes actions given state observations from an environment; the environment subsequently emits rewards and new state observations, based on the agent's actions
- **Technical Direction**
  - Job allocation for Dragonfly Networks
  - Adaptive P-State control for HPC workloads
  - Multi-level Memory Management
- **Results/Impact**
  - Developed RL MODL (Mathematical Optimizations for Deep Learning) library
  - Quantization study of RL showing results on continuous control problems

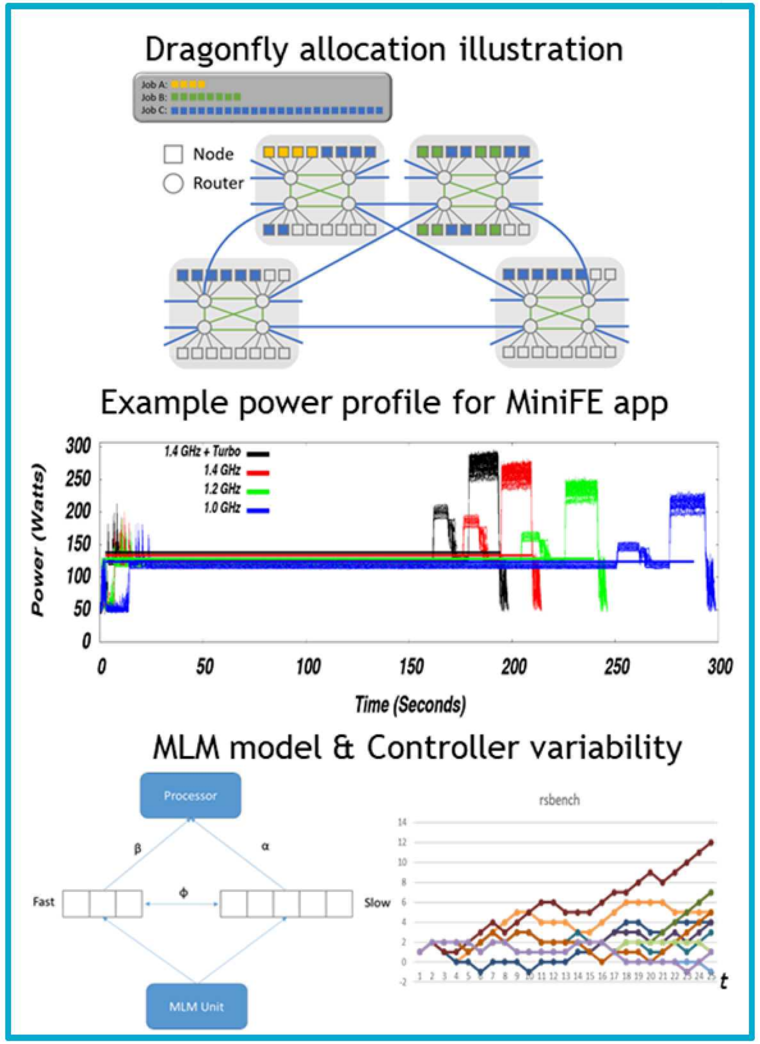

Main Focus Areas: Reduced computational cost → Maintain International HPC/HPDA leadership → Reduce the design cycle time; Increased production throughput; Re-think the surveillance program for the 21st century

Parallel Training of Deep Residual Neural Networks

- PI: Eric Cyr, ASCR funded project

Learning Hardware

- **Problem**
  - Deep Learning is often viewed as a black box that requires large training sets and significant training time
  - High-consequence decisions made with ML analysis need to be explainable and credible
- **Technical Approach**
  - Training a Residual Neural Network (ResNet) is cast as an optimal control problem subject to nonlinear dynamics
  - Classical forward and backward propagation through network layers are replaced by a parallel MultiGrid Reduction In Time (MGRIT) iteration in the layer domain.
- **Results/Accomplishments**
  - Unique "Layer-Parallel" approach provides scalable speed-up over serial stochastic gradient descent approach to training (~16x at 2048 layers)
  - Theoretical basis to learning provided by dynamic system optimization methodology

Layer-Parallel training for ResNet

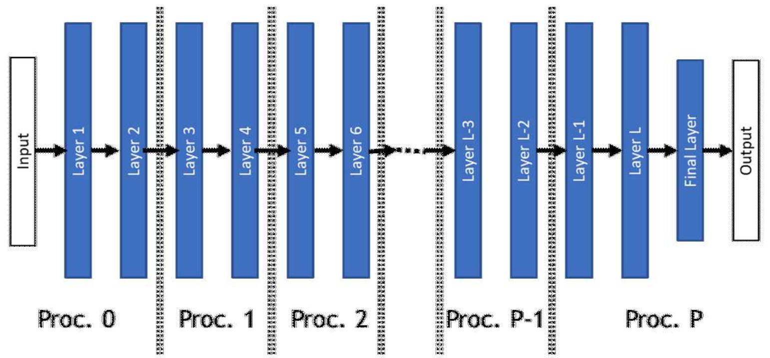

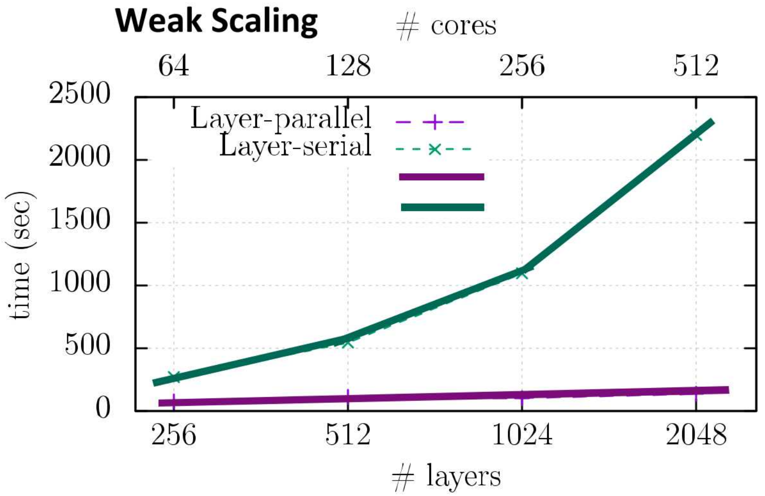

S. Guenther, L. Ruthotto, J. B. Schroder, E. C. Cyr, N. R. Gauger, Layer-Parallel Training of Deep Residual Neural Networks, preprint arXiv:1812.04352, 2018.

Main Focus Areas: Reduced computational cost

Maintain International HPC/HPDA leadership

Increased production throughput

Reduce the design cycle time

Re-think the surveillance program for the 21st century

Conclusions

- Machine learning will provide new capabilities for scientific and engineering applications
  - Reduced order surrogate models for scientific/engineering problems
    - Could help us learn what is wrong/missing in physics models and aid in experimental design
  - Ability to identify anomalies and regions of interest in inspection and surveillance data
  - Correlating and certifying simulation and experimental results
  - For agility of application workflows (automating processes)

- Machine learning will provide new capabilities for HPC system administrators, facilities, dev-ops, and system software
  - Help model complex behaviors (e.g., failures, degradation, energy)
  - Automate/adapt usage to comply with more complex policy (e.g., energy consumption)
  - Adaptable resource management (e.g., network, memory, storage, energy)
  - “Smart” data-movement for Exascale runtimes

Questions?