MATLAB EXPO 2017 KOREA

4월 27일, 서울

등록 하기 matlabexpo.co.kr

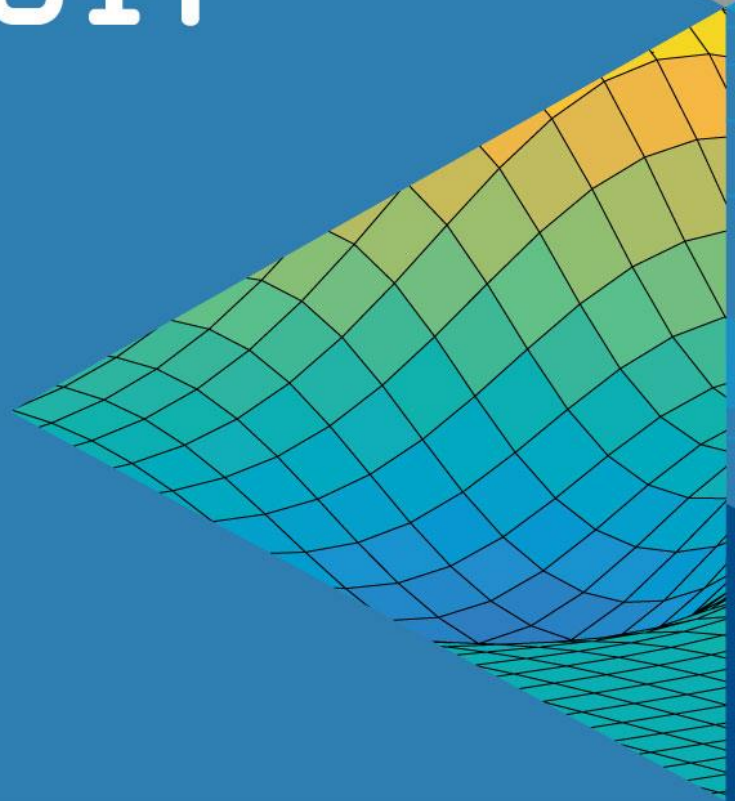

Parallel Computing with MATLAB and Simulink

정승혁 과장
Senior Application Engineer
MathWorks Korea

Who Uses Parallel Computing?

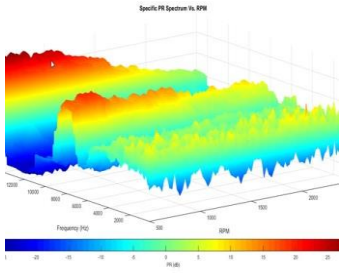

Automotive Test Data Analysis and Visualization
Validation time reduced by 40-50%
3-4 months of development time saved

Heart Transplant Studies
4 weeks reduced to 5 days
6X speedup in process time

Design and Build Wave Energy Farm
Sensitivity studies accelerated 12x

Discrete-Event Model of Fleet Performance
Simulation time reduced from months to hours
20X faster simulation time
Linkage with Neural Network Toolbox

Calculating Derived Market Data
Implementation time reduced by months
Updates loaded 8X faster

Overcome Challenges With MathWorks Parallel Computing Tools

- Parallel Computing-key challenges
  - Need faster insight to bring competitive products to market quickly
  - Size and complexity of analytical problems is growing across industries
- Key Takeaways
  - Save engineering and research time and focus on results
  - Leverage computational power of broadly available hardware

    (Multicore Desktops, GPUs, Clusters)
  - Seamlessly scale from your desktop to clusters or the cloud

Agenda

- Parallel computing in MATLAB and Simulink
- Accelerate applications with NVIDIA GPUs
- Scaling to clusters and clouds
- Summary

Parallel Computing Paradigm

Multicore Desktops

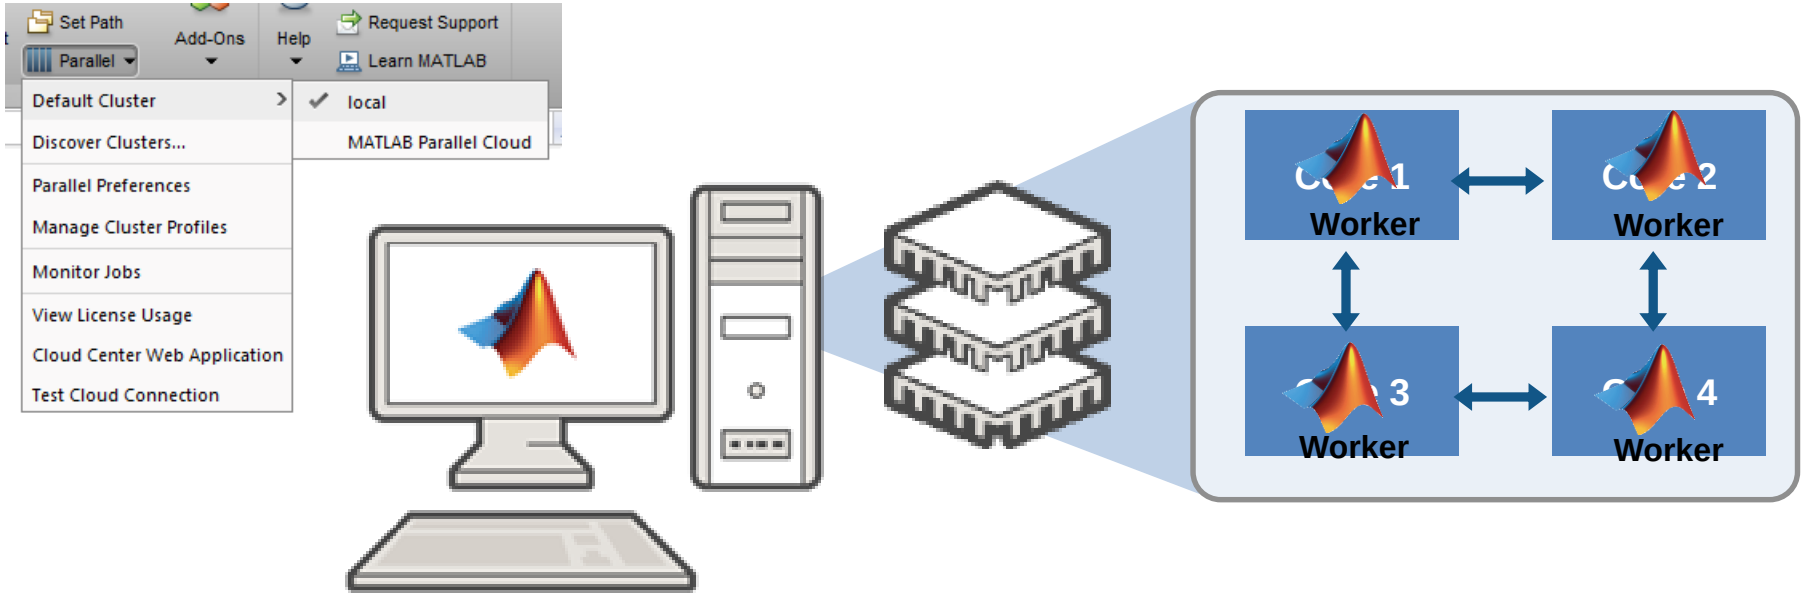

[MATLAB multicore]

Parallel Computing Paradigm

Clusters

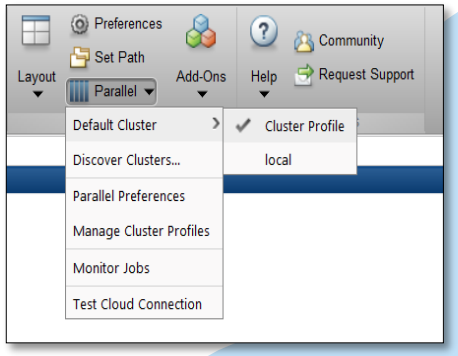

Parallel-enabled Toolboxes (MATLAB® Product Family)

Enable parallel computing support by setting a flag or preference

Image Processing

Batch Image Processor, Block Processing, GPU-enabled functions

Statistics and Machine Learning

Resampling Methods, k-Means clustering, GPU-enabled functions

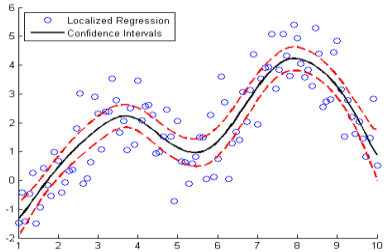

Neural Networks

Deep Learning, Neural Network training and simulation

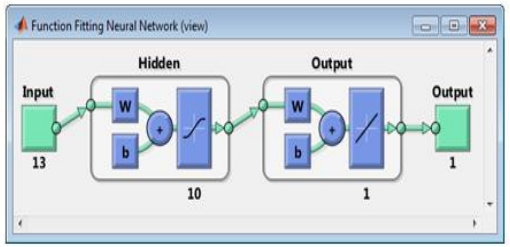

Signal Processing and Communications

GPU-enabled FFT filtering, cross correlation, BER

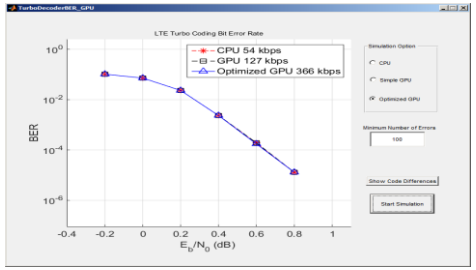

Computer Vision

Parallel-enabled functions in bag-of-words workflow

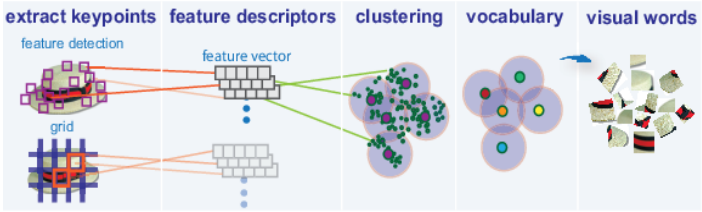

Optimization

Parallel estimation of gradients

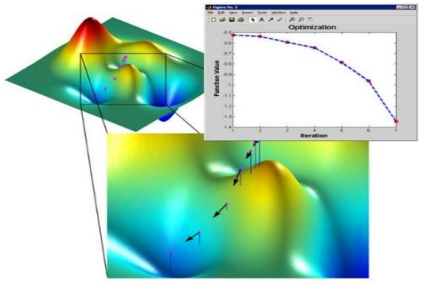

[Other Parallel-enabled Toolboxes]

Classification learner : Enable Parallel

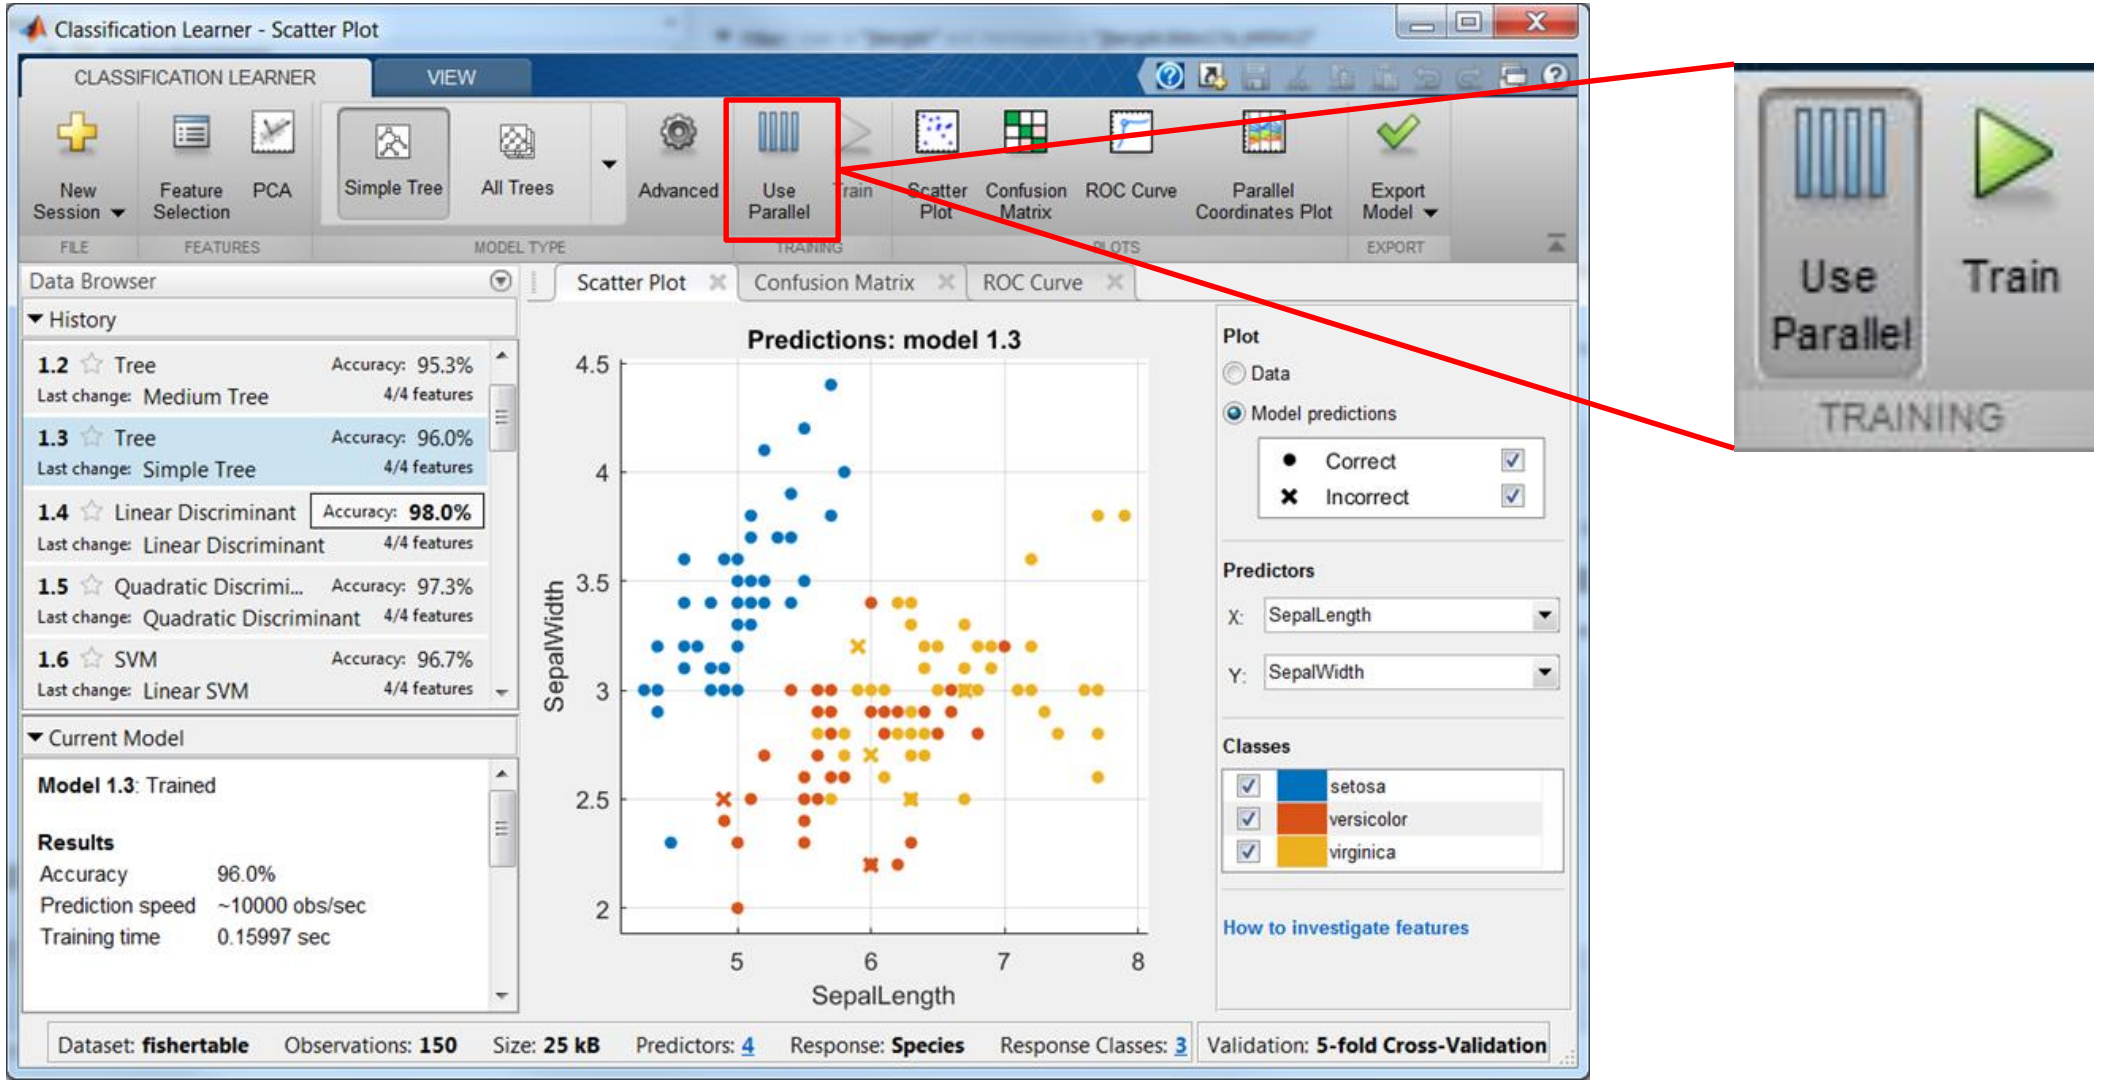

Parallel-enabled Toolboxes (Simulink® Product Family)

Enable parallel computing support by setting a flag or preference

Simulink Design Optimization

Response optimization, sensitivity analysis, parameter estimation

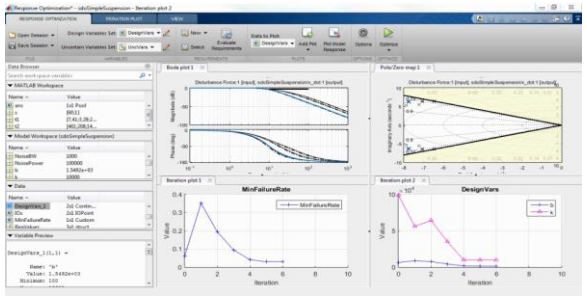

Simulink Control Design

Frequency response estimation

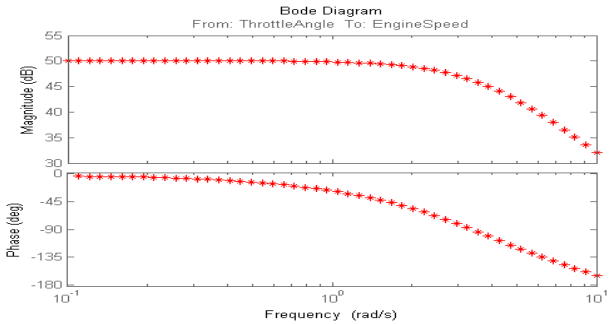

Communication Systems Toolbox

GPU-based System objects for Simulation Acceleration

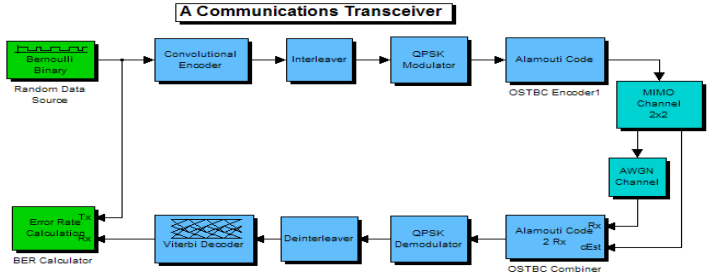

Simulink/Embedded Coder

Generating and building code

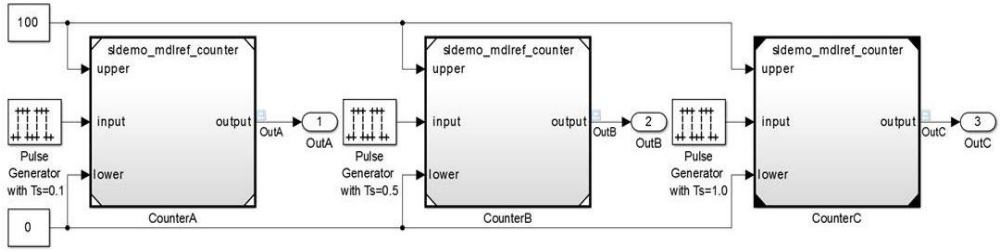

[Other Parallel-enabled Toolboxes]

Explicit Parallelism: Independent Tasks or Iterations

Simple programming constructs using `parfor, parfeval`

- Examples: parameter sweeps, Monte Carlo simulations
- No dependencies or communications between tasks

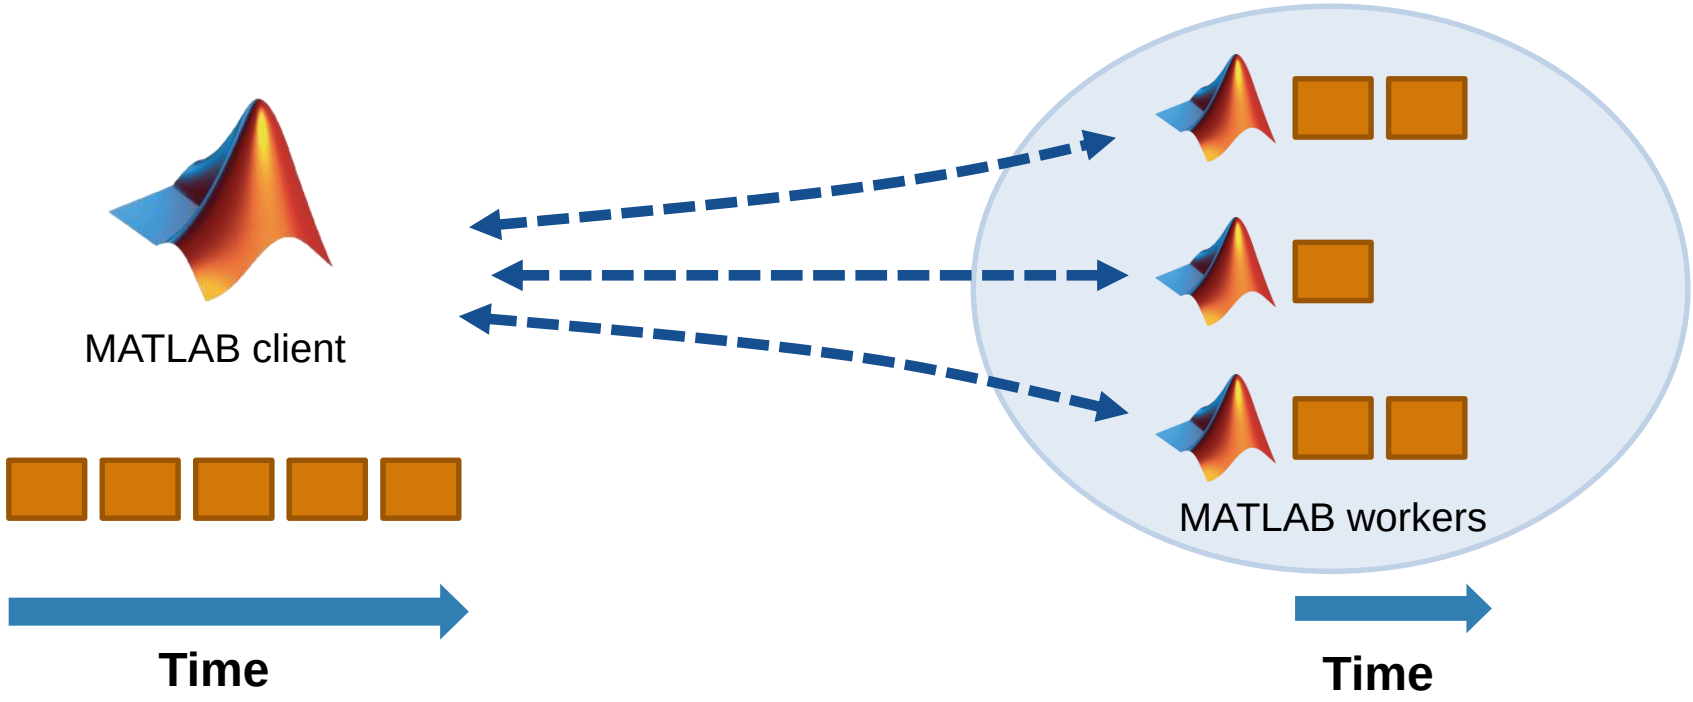

Energy Production – World's First Operating Wave Farm

Carnegie Wave Energy

Goal: Develop unique technology for generating electric power from ocean waves

A CETO unit ready for deployment in the wave farm.

Learn More

Challenges

- Analyze loads and estimate energy output without building scale model of entire system
- Run simulations for a range of configurations, sea conditions and faults

Why Parallel Computing

- Sensitivity studies accelerated

*"Our sensitivity studies require numerous simulations because we typically simulate 15 to 20 sea states for each parameter value we vary. With Parallel Computing Toolbox we can run **simulations in parallel**, and with a twelve-core computer we see an almost **twelvefold increase in speed**."*

Jonathan Fiévez, Carnegie Wave Energy

Leverage Parallel Computing with Simulink

- Reduce the total amount of time it takes to…

•Run multiple independent simulations (E.g. Parameter sweeps, Monte Carlo Analysis)

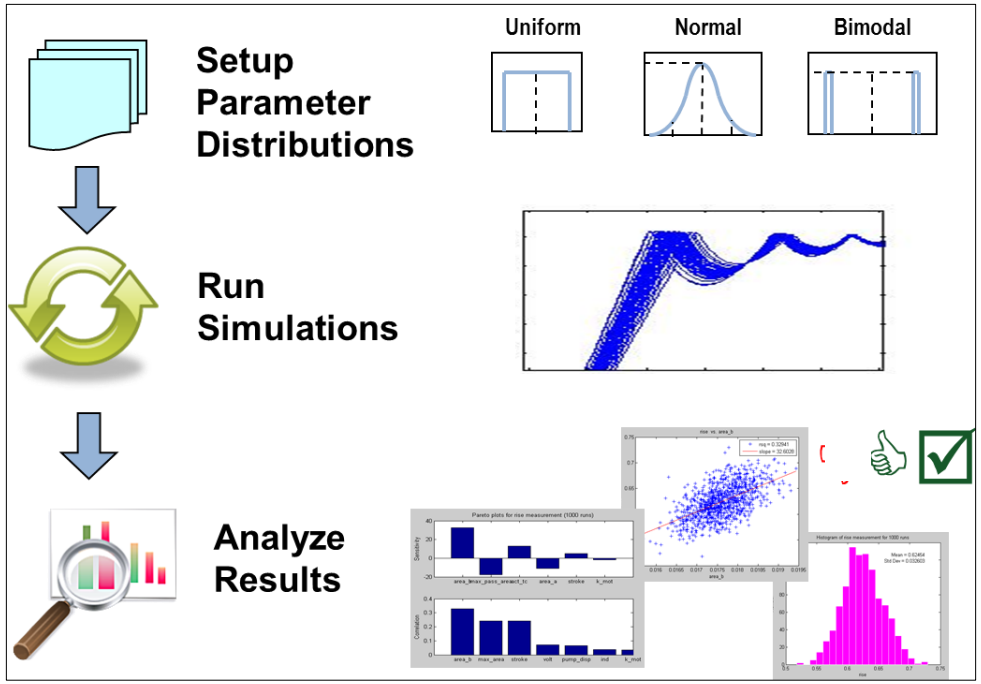

Update models containing large model reference hierarchies

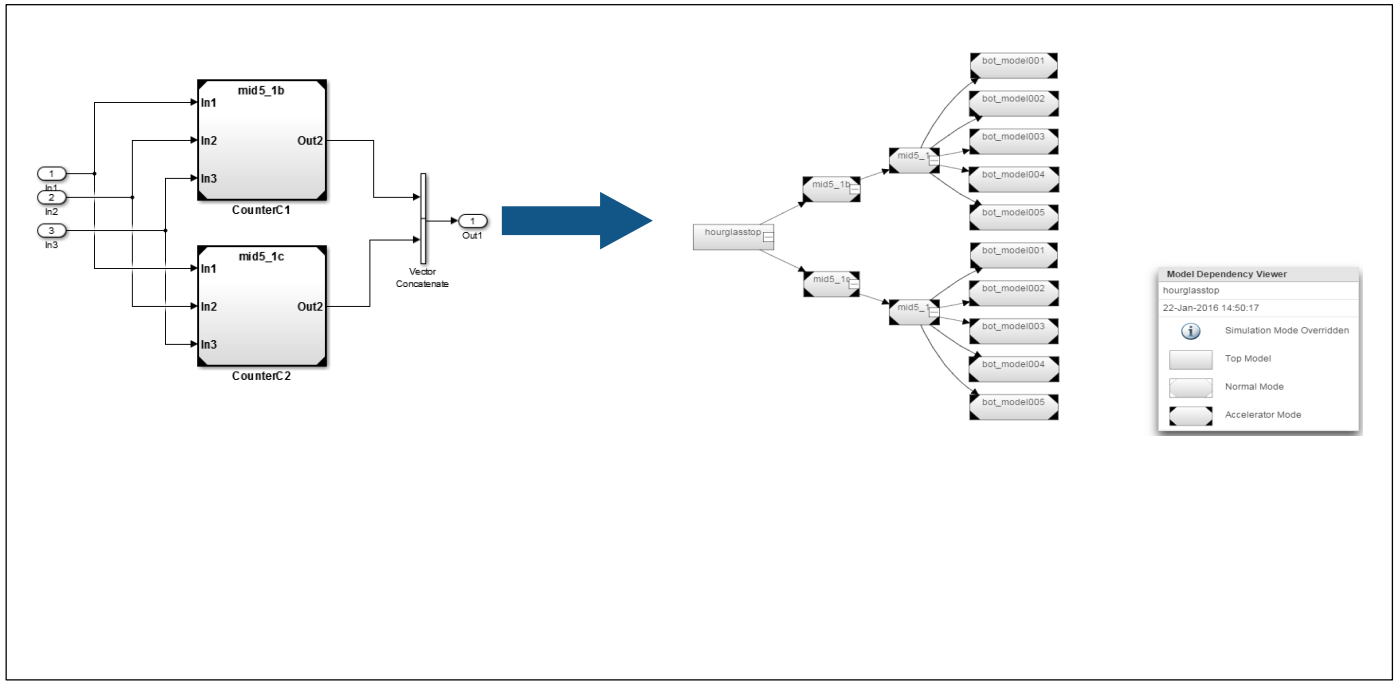

Easy execution of parallel simulations using `Parsim`

Directly run multiple parallel simulations from the `parsim` command

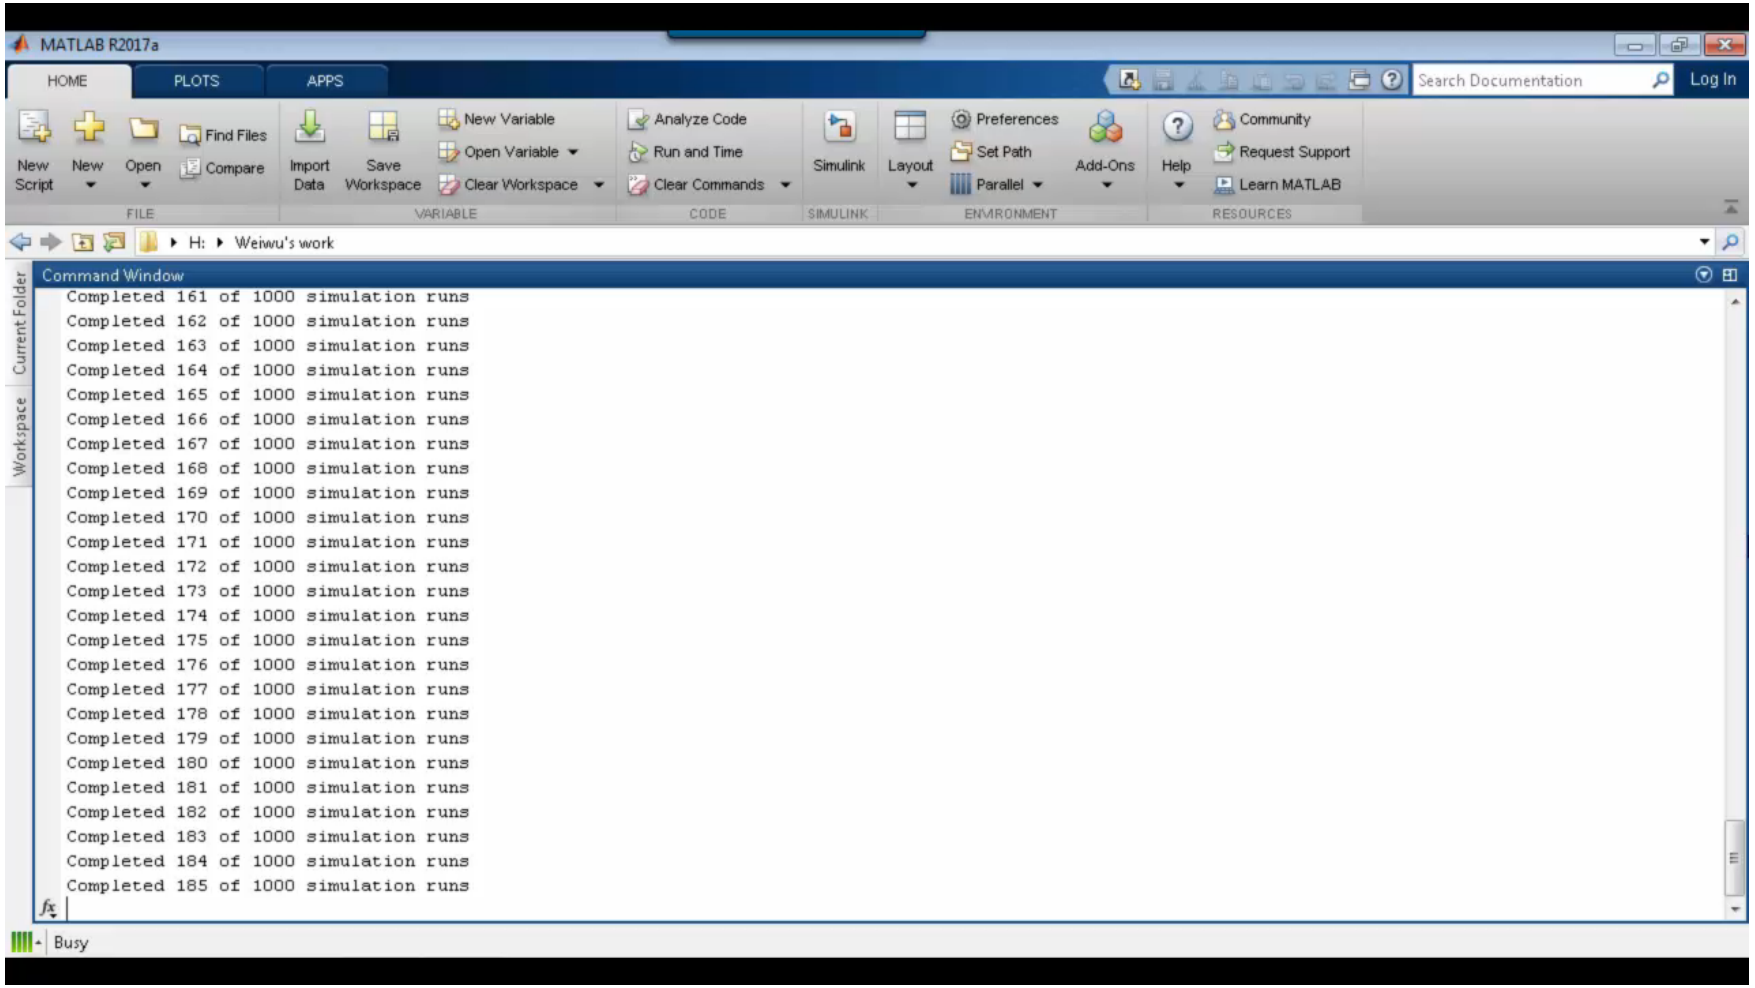

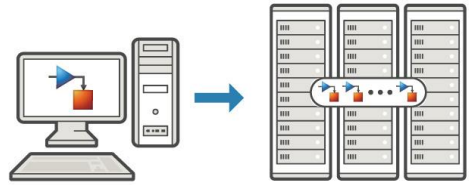

- Enables customers to easily use Simulink with parallel computing
- Simplifies customers' large simulation runs and improves their productivity
- We manage the parallel setup so customers can focus on their simulations

Agenda

- Parallel computing in MATLAB and Simulink
- Accelerate applications with NVIDIA GPUs
- Scaling to clusters and clouds
- Summary

Parallel Computing Paradigm

Going Parallel: GPUs

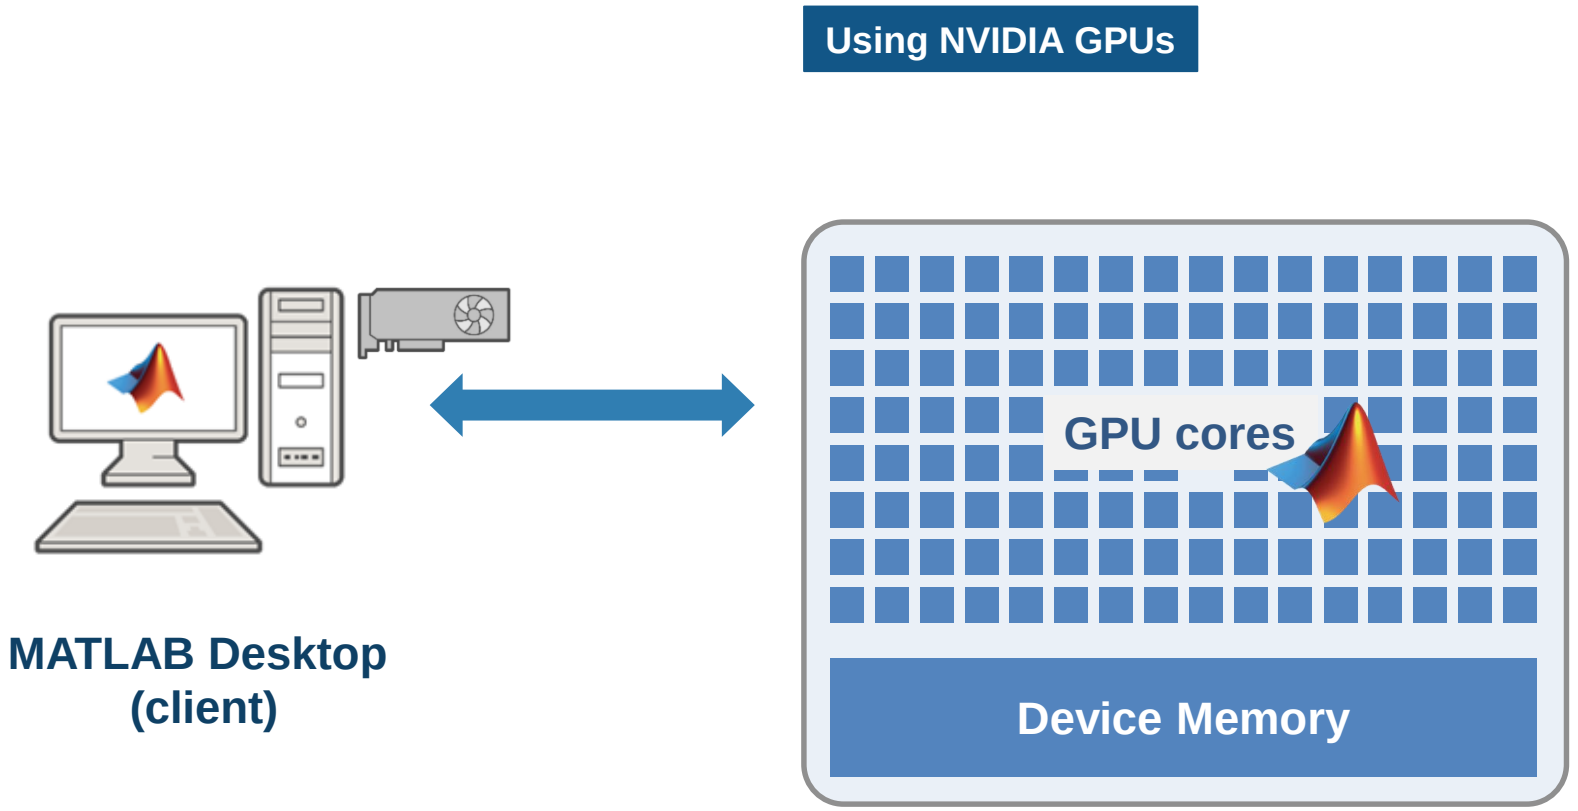

Speed-up using NVIDIA GPUs

- Ideal Problems
  - Massively Parallel and/or Vectorized operations
  - Computationally Intensive
  - Algorithm consists of supported functions
- 300+ GPU-enabled MATLAB functions
- Additional GPU-enabled Toolboxes
  - Neural Networks
  - Image Processing
  - Communications
  - Signal Processing

..... Learn More

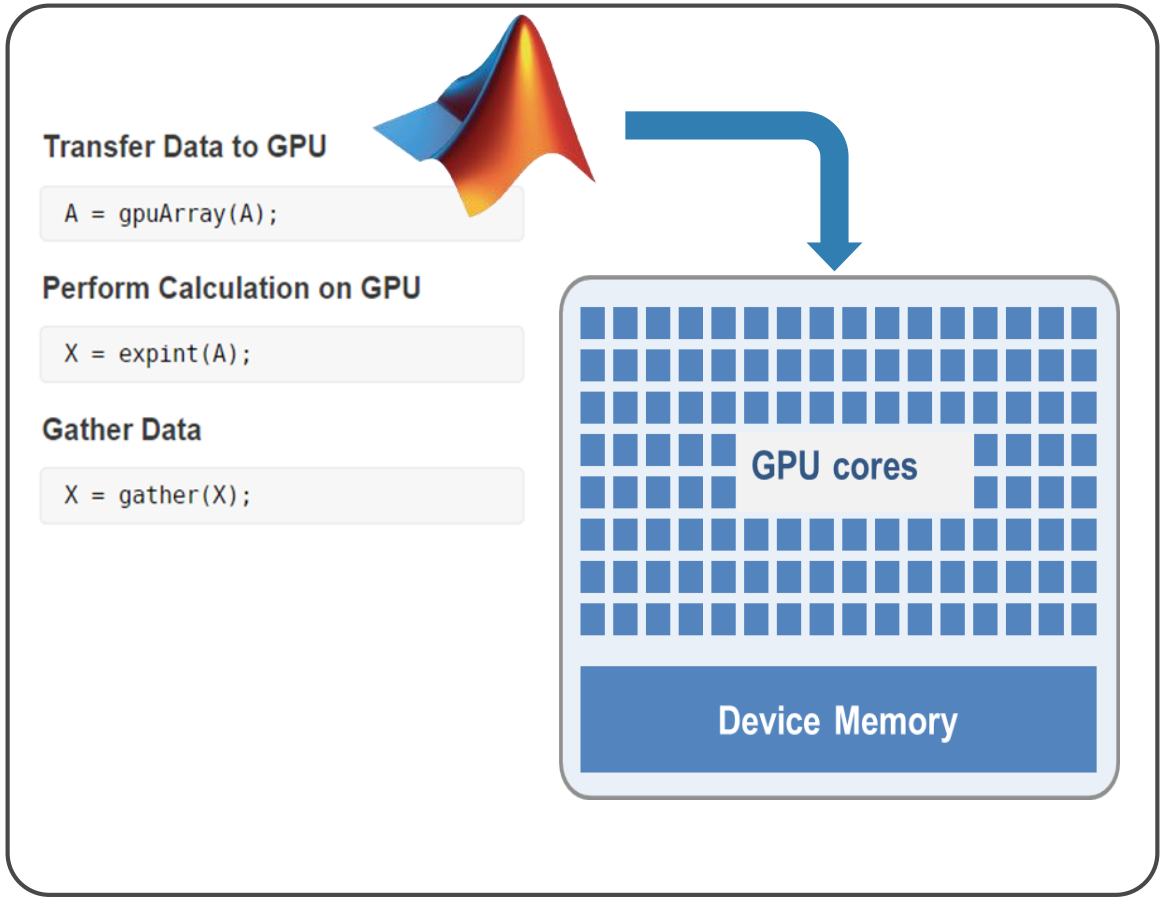

Signal Processing – Acoustic Data Analysis

NASA Langley Research

Goal: Accelerate the analysis of sound recordings from wind tunnel tests of aircraft components

Learn More

Challenges

- Legacy code took 40 mins to analyze single wind tunnel test data
- Reduce processing time to make on-the-fly decisions and identify hard ware problems

Why GPU Computing

- Computations completed 40 times faster.

*"Many operations we perform, including FFTs and matrix multiplication, are **GPU-enabled MATLAB functions** .Once we developed the initial MATLAB code for CPU execution, it took 30 minutes to get our algorithm working on the GPU—**no low-level CUDA programming** was needed. The addition of GPU computing with Parallel Computing Toolbox cut it to **under a minute**, with most of that time spent on data transfer"*

Christopher Bahr, NASA

Agenda

- Parallel computing in MATLAB and Simulink
- Accelerate applications with NVIDIA GPUs
- Scaling to clusters and clouds
- Summary

Why parallel computing matters

Scaling case study with a compute cluster

$$M\ddot{x} + C\dot{x} + Kx = F$$

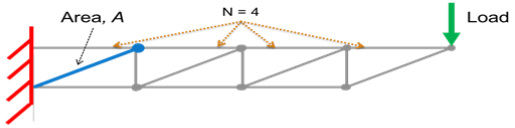

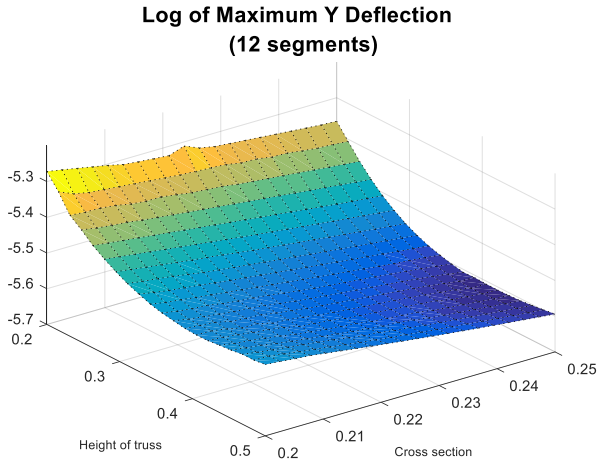

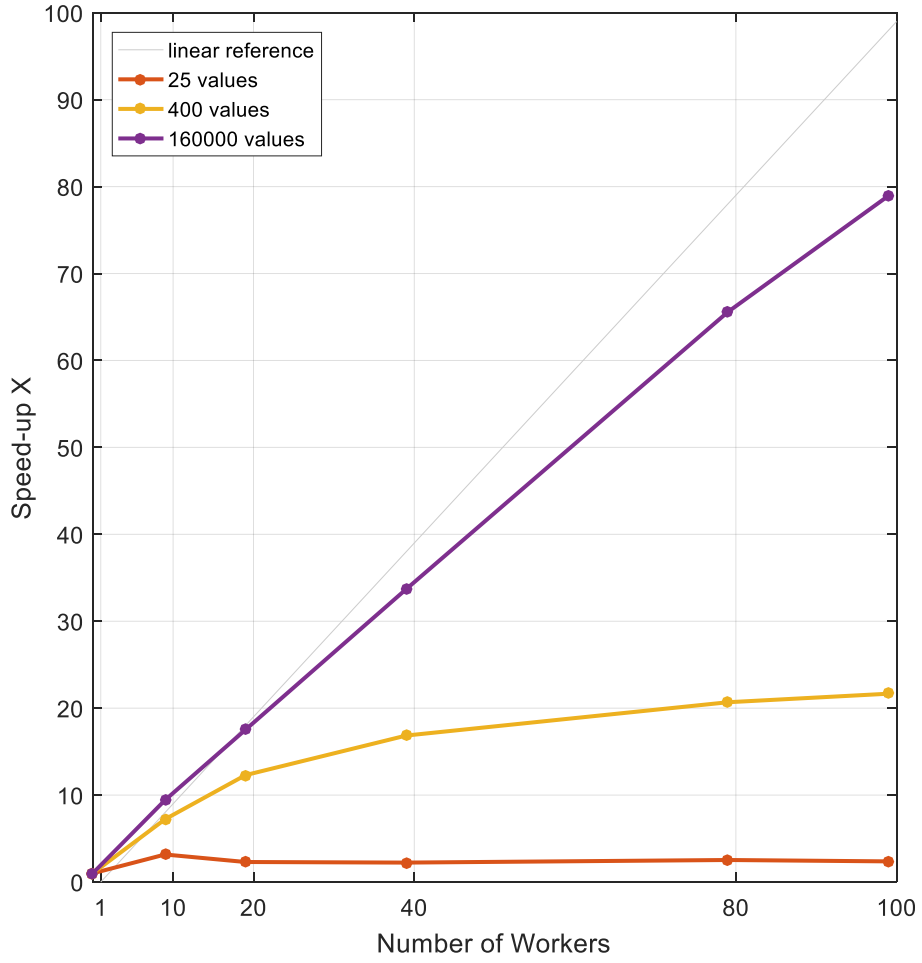

Workers in pool	Compute time (minutes)		
	160e3 values	400 values	25 values
1	140	0.38	0.03
10	15	0.05	0.01
20	8.0	0.03	0.01
40	4.2	0.02	0.01
80	2.1	0.02	0.01
100	1.8	0.02	0.01

Processor: Intel Xeon E5-class v2
16 physical cores per node
MATLAB R2016a

Scaling to a computer cluster or cloud

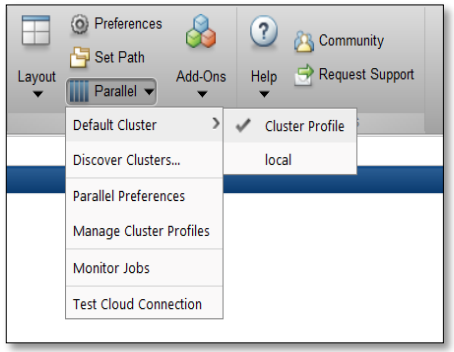

MATLAB Desktop (client)

Microsoft Azure

amazon web services™

Cluster of computers

Worker

Implement Machine Learning Portfolio Allocation Models in the Cloud

Aberdeen Asset Management

Goal: Improve asset allocation strategies by creating model portfolios with machine learning techniques

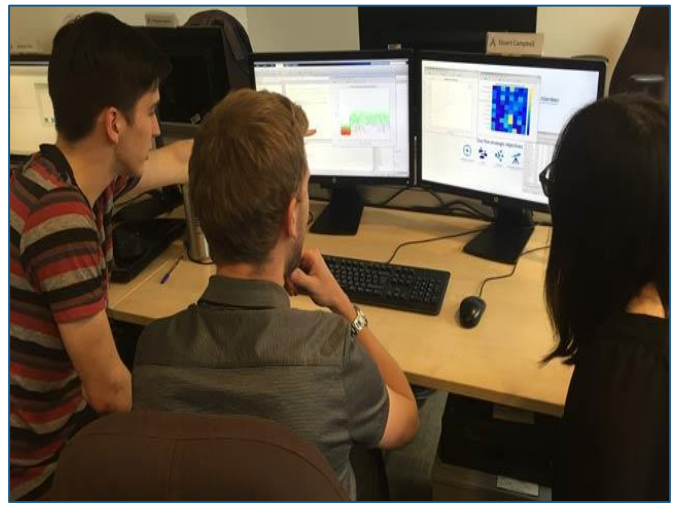

Challenges:

- Use MATLAB to develop classification tree, neural network, and support vector machine models,

Why Parallel Computing:

- Use MATLAB Distributed Computing Server to run models in the cloud
  - Back tested the trained models on more than 15 years of historical data
  - Run the parallel execution on an onsite cluster with 80 workers
  - Redeployed models to 80 workers running on Microsoft Azure virtual machines
- *Portfolio performance goals supported*
- *Processing times cut from 24 hours to 3*
- Multiple types of data easily accessed

"The widespread use of MATLAB in the finance community is a real advantage. Many university students learn MATLAB and can contribute right away when they join our team during internship programs. In addition, the strong MATLAB libraries developed by academic researchers help us explore all the possibilities of this programming language."

Emilio Llorente-Cano, Aberdeen Asset Management

What's new in 16b and 17a?

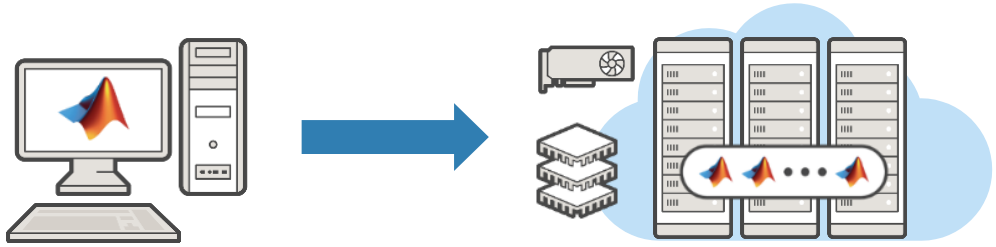

R2016b

- `tall` array support for big data
- Measure data sent to workers using `ticBytes` and `tocBytes`
- Cloud offerings with K80-equipped GPUs

R2017a

- Simplified parallel Simulink simulations using `parsim`
- Send data to client using `DataQueue` and `PollableDataQueue`
- Train a single deep learning network with multiple CPUs or multiple GPUs


MathWorks®