MATLAB EXPO 2017

Parallel Computing with MATLAB and Simulink

Alka Nair
Application Engineer
MathWorks India Private Limited

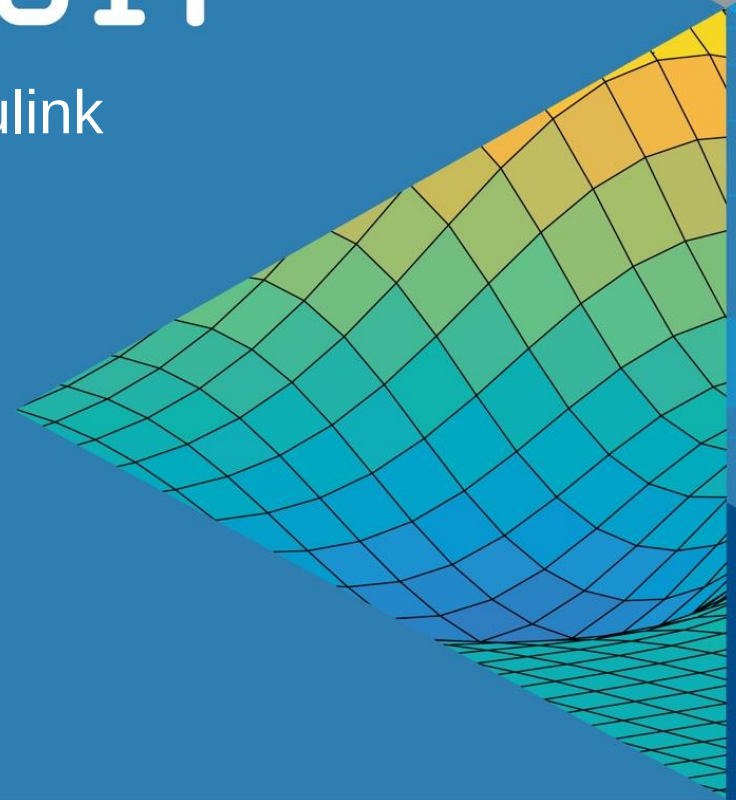

Why Parallel Computing ?

- Size and complexity of analytical problems is growing across industries
- Need faster insight to bring competitive products to market quickly
- Hardware is becoming powerful: Leverage computational power of multicore desktops, GPUs, clusters

Key Takeaways

- Overcome challenges with MathWorks Parallel Computing Tools
  - Save engineering and research time and focus on results
  - Leverage computational power of broadly available hardware with minimal changes to your existing code
    - (Multicore Desktops, GPUs, Clusters)
  - Seamlessly scale from your desktop to clusters or the cloud
  - Speed-up analysis of Big Data using built-in parallel computing capabilities

Where is Parallel Computing Used?

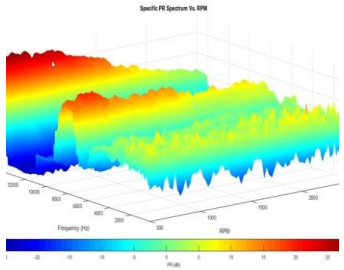

Bosch Develops Platform for Automotive Test Data Analysis and Visualization
Validation time **reduced by 40-50%**
3-4 months of development time saved

Heart Transplant Studies
4 weeks reduced to 5 days

Carnegie Wave Energy Designs and Builds Wave Energy Farm
Sensitivity studies accelerated 12x

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

Commerzbank Develops Production Software Systems for Calculating Derived Market Data
Implementation time **reduced by months**
Updates loaded 8X faster

Agenda

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

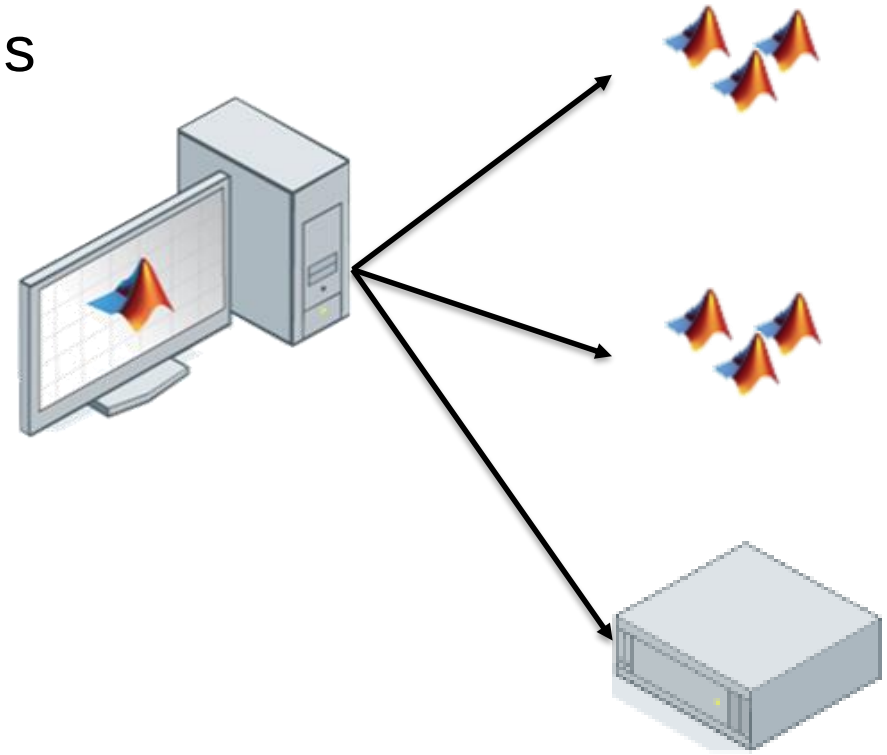

Classification learner demo

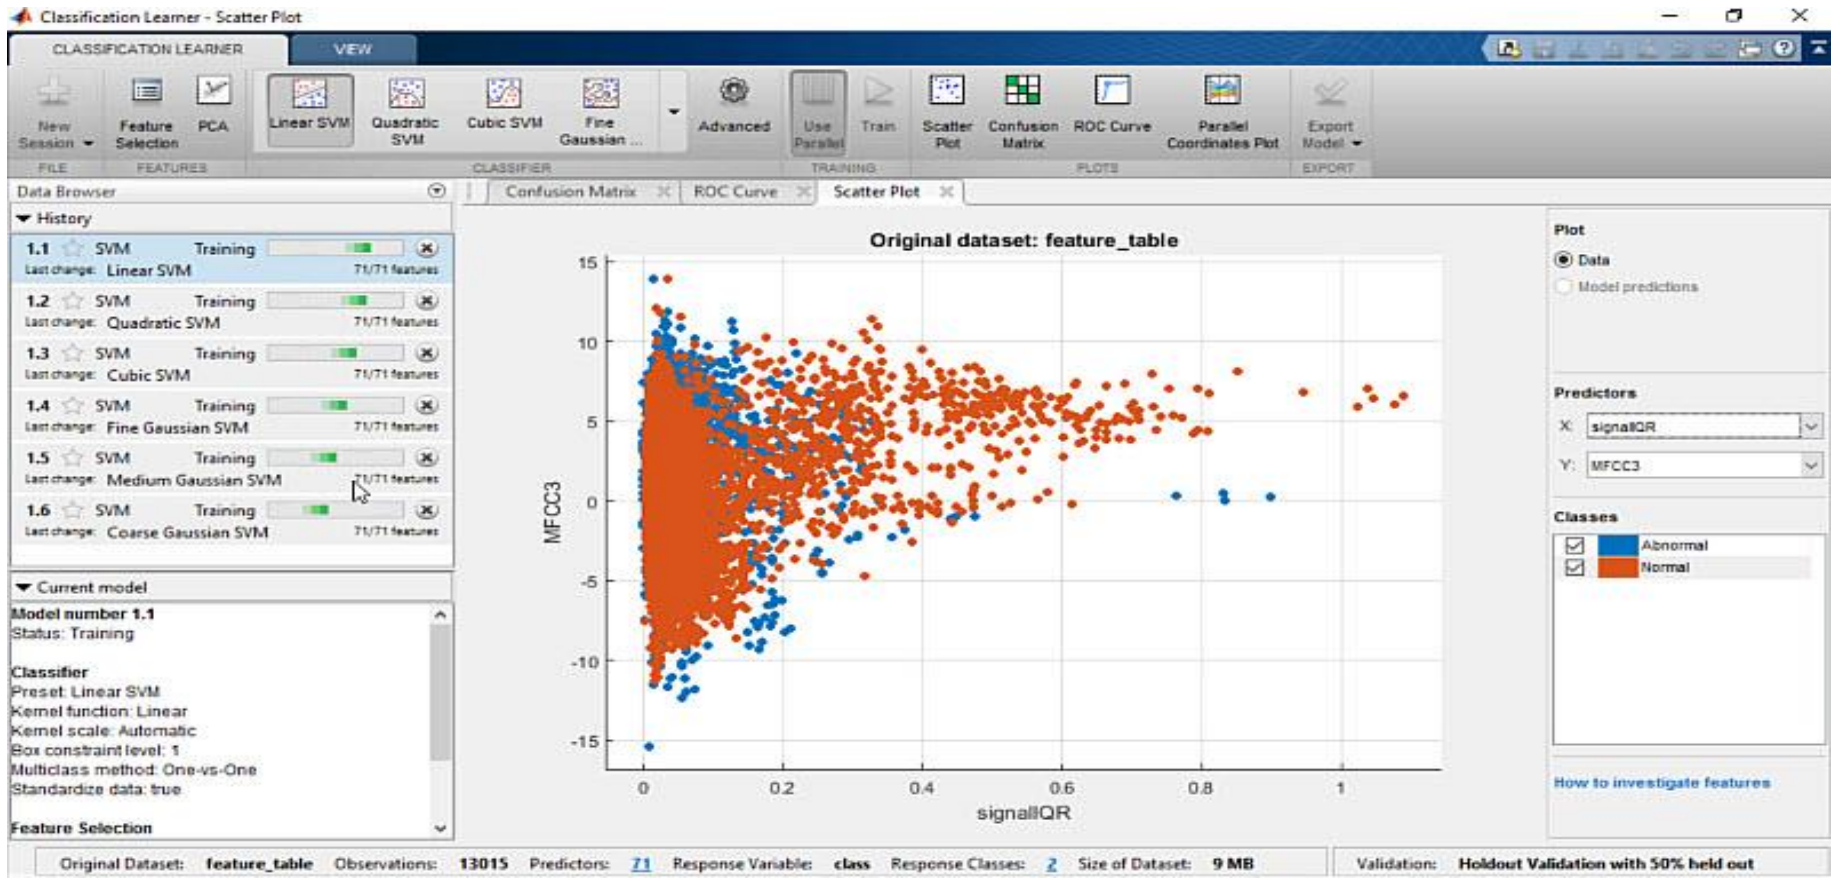

Parallel Computing Paradigm - Hardware

Multicore Desktops

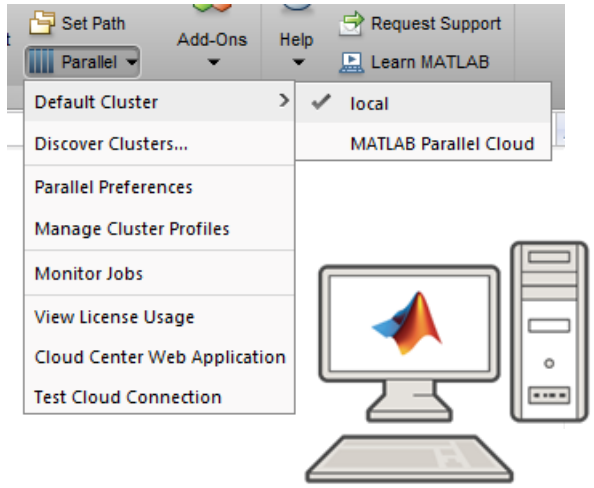

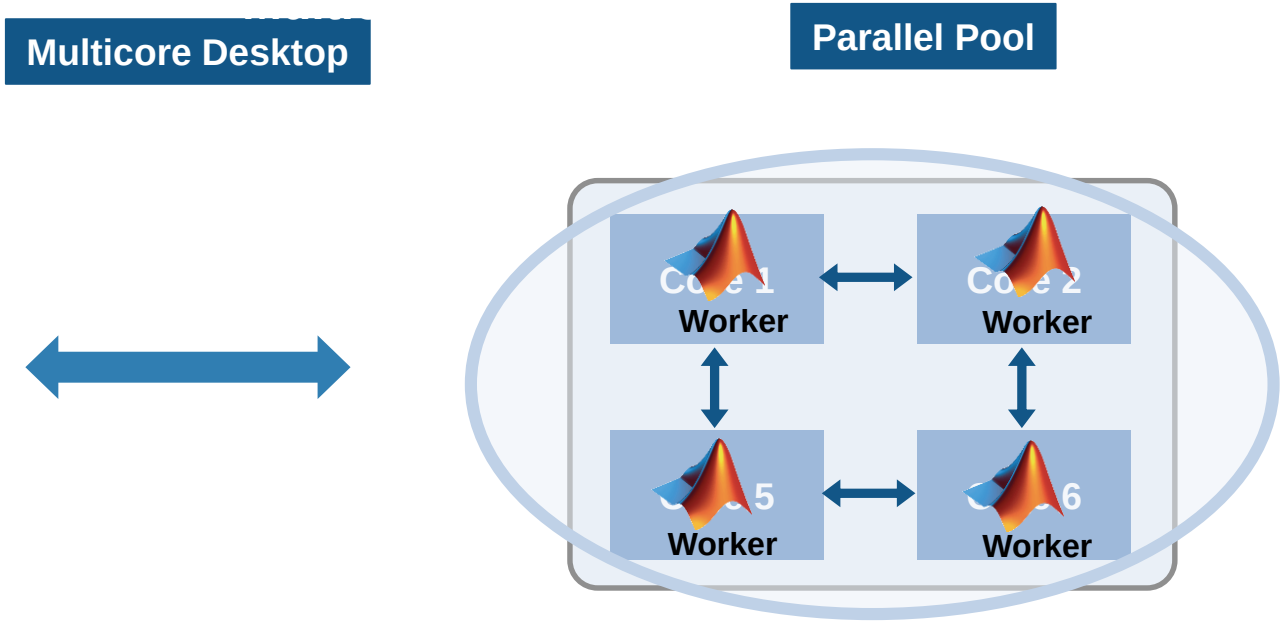

Parallel Computing Paradigm - Hardware

Cluster Hardware

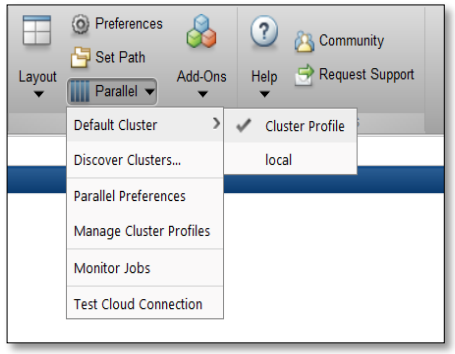

MATLAB Desktop
(client)

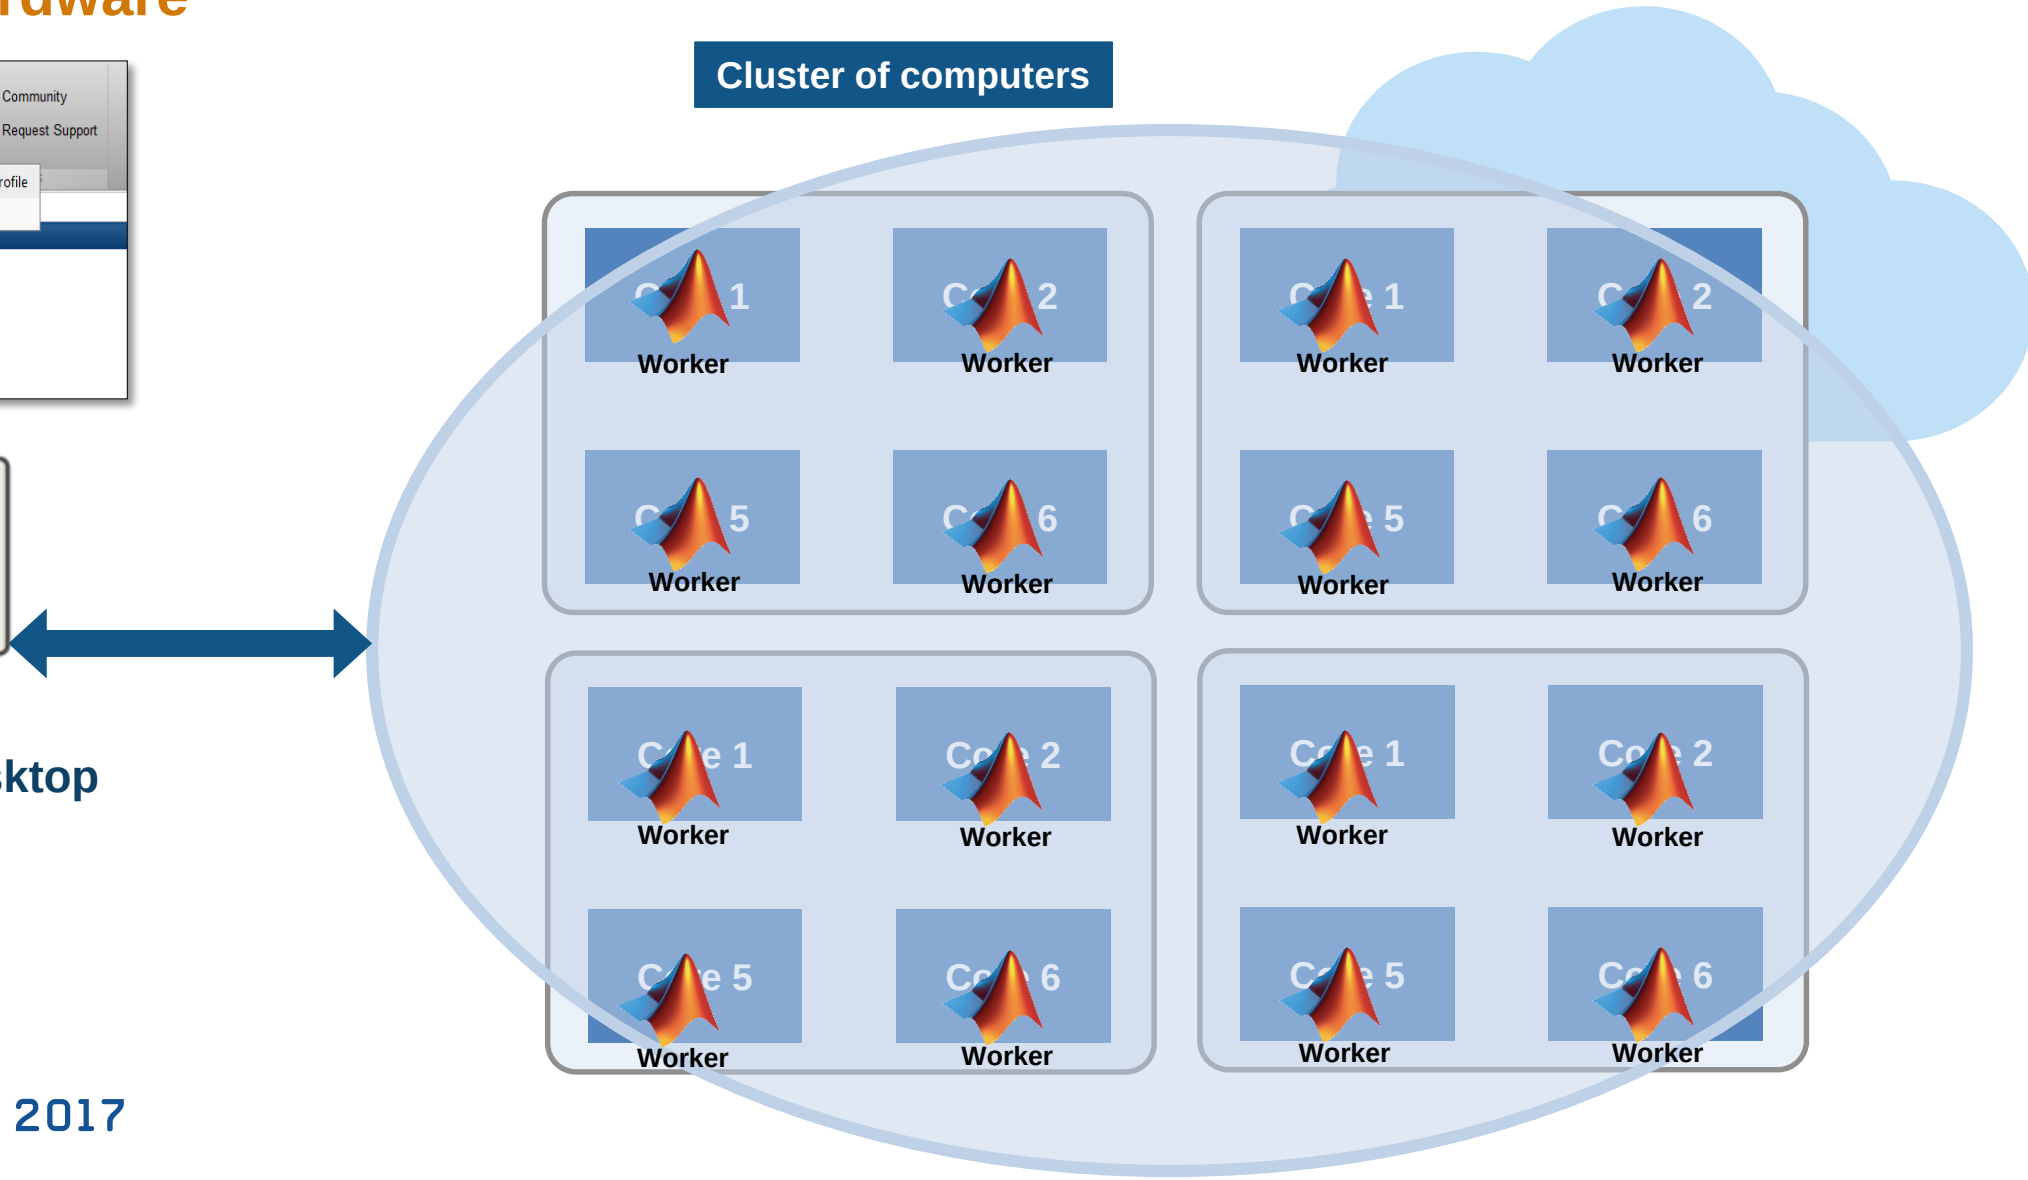

Programming Parallel Applications – Programming Constructs

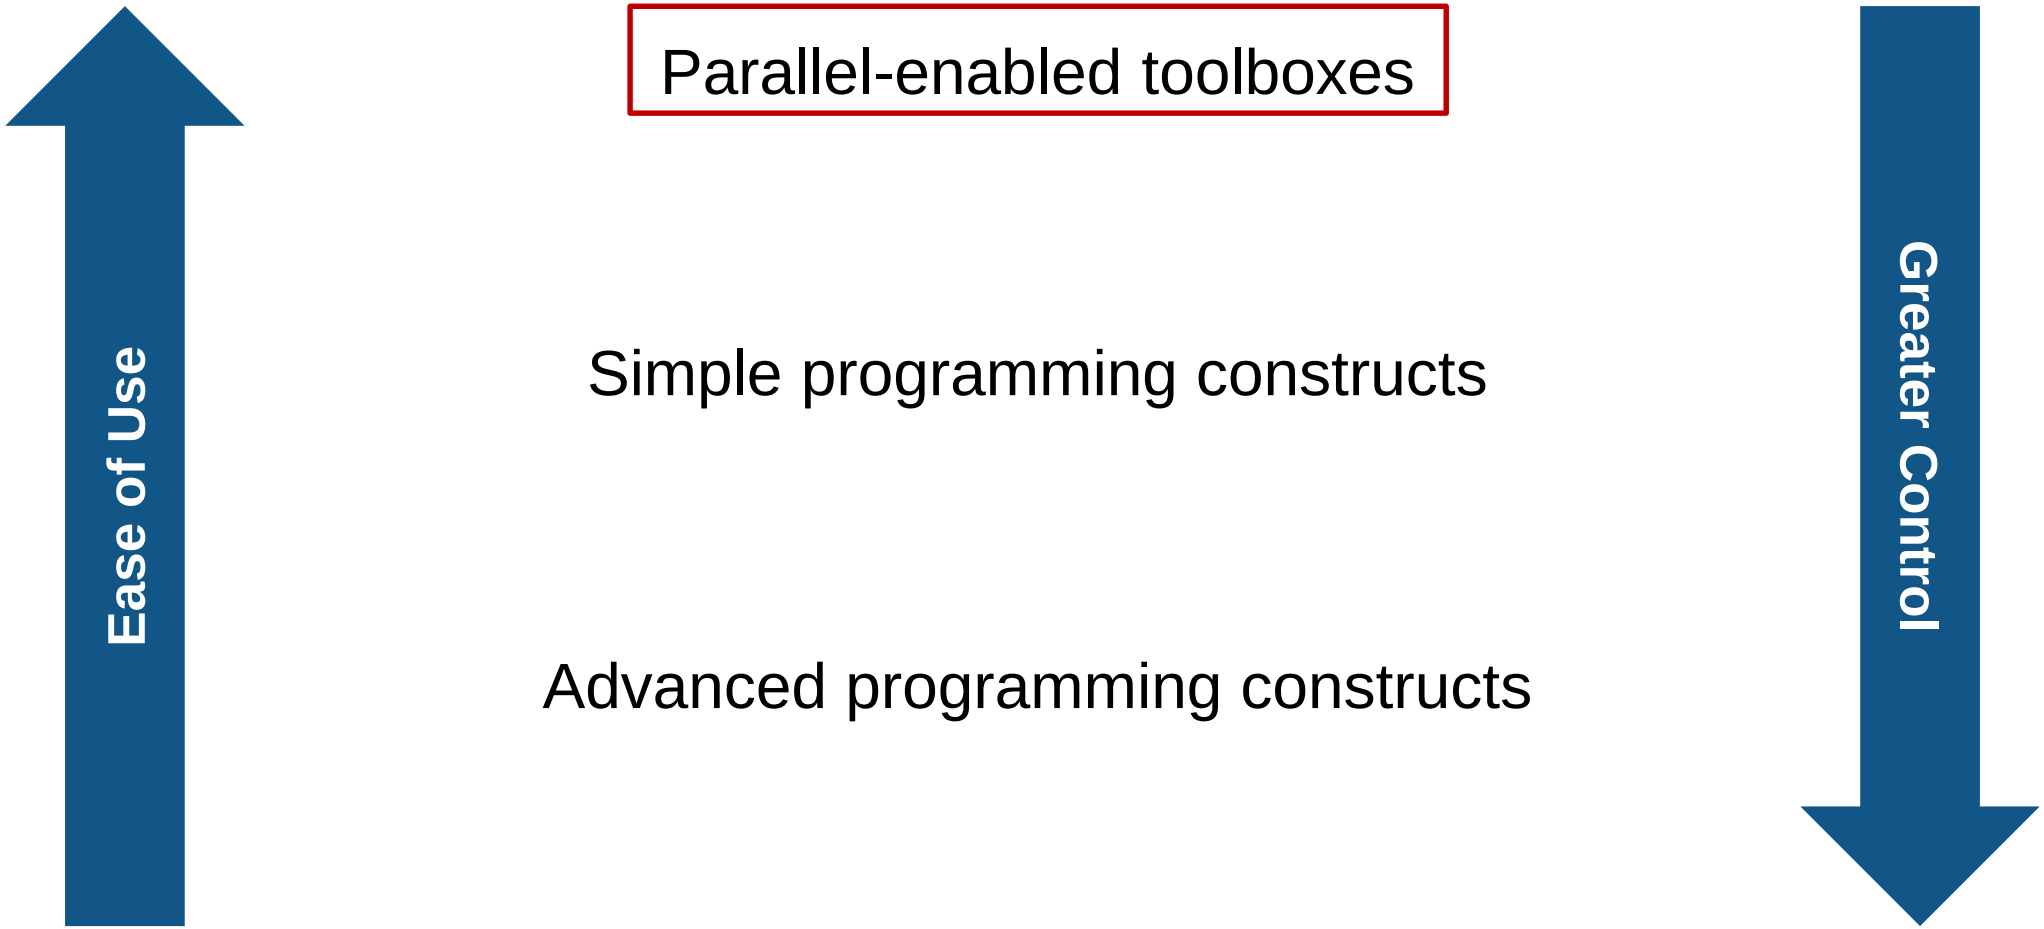

Parallel Computing: Neural Network Toolbox

```
%% Load data set
[x, t] = bodyfat_dataset;
%% Define the network
net1 = feedforwardnet(10);
%% Use parallel Computing to train the Network
net2 = train(net1,x,t,'useParallel','yes');
y = net2(x,'useParallel','yes');
```

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

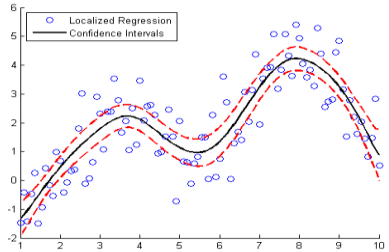

Neural Networks

Deep Learning, Neural Network training and simulation

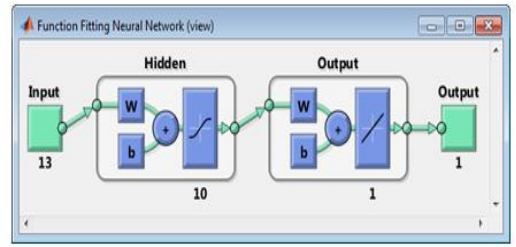

Signal Processing and Communications

GPU-enabled FFT filtering, cross correlation, BER

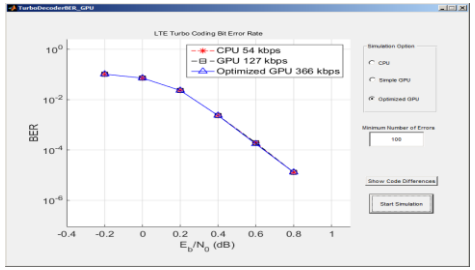

Computer Vision

Parallel-enabled functions in bag-of-words workflow

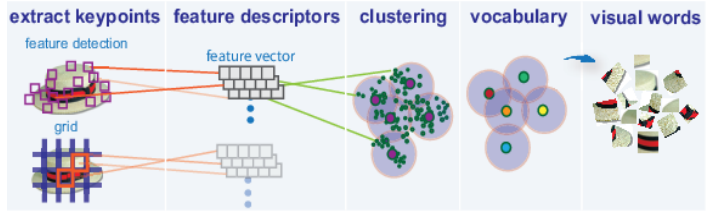

Optimization

Parallel estimation of gradients

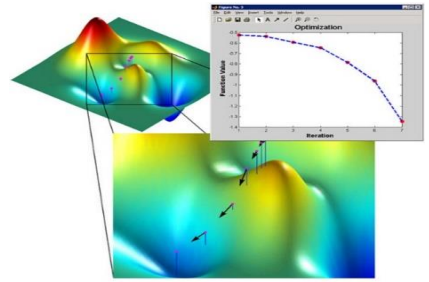

Other Parallel-enabled Toolboxes

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

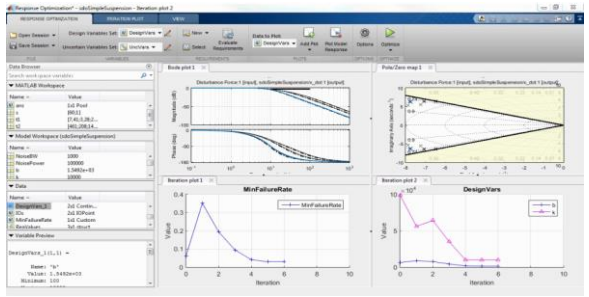

Simulink Control Design

Frequency response estimation

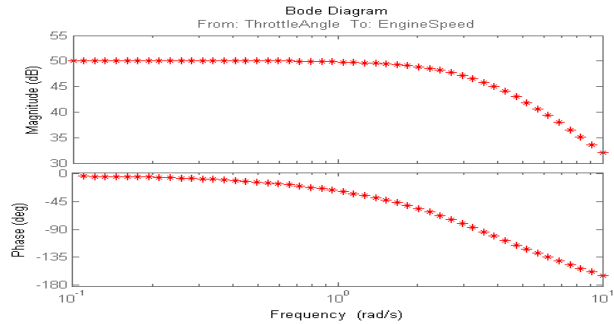

Communication Systems Toolbox

GPU-based System objects for Simulation Acceleration

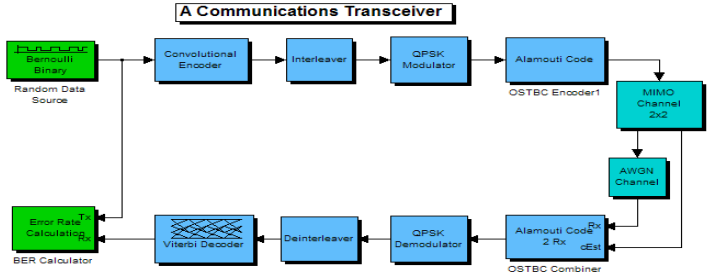

Simulink/Embedded Coder

Generating and building code

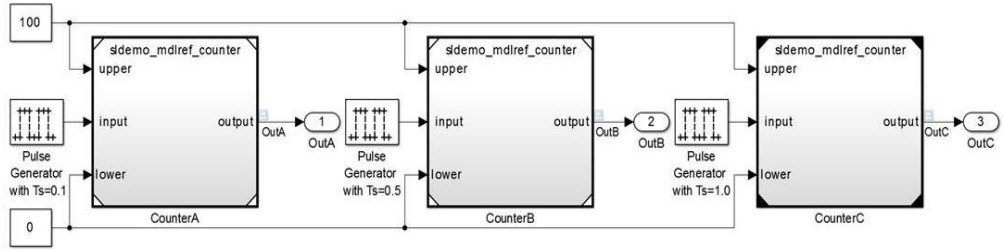

[Other Parallel-enabled Toolboxes]

Data Intensive: Big Data support in MATLAB

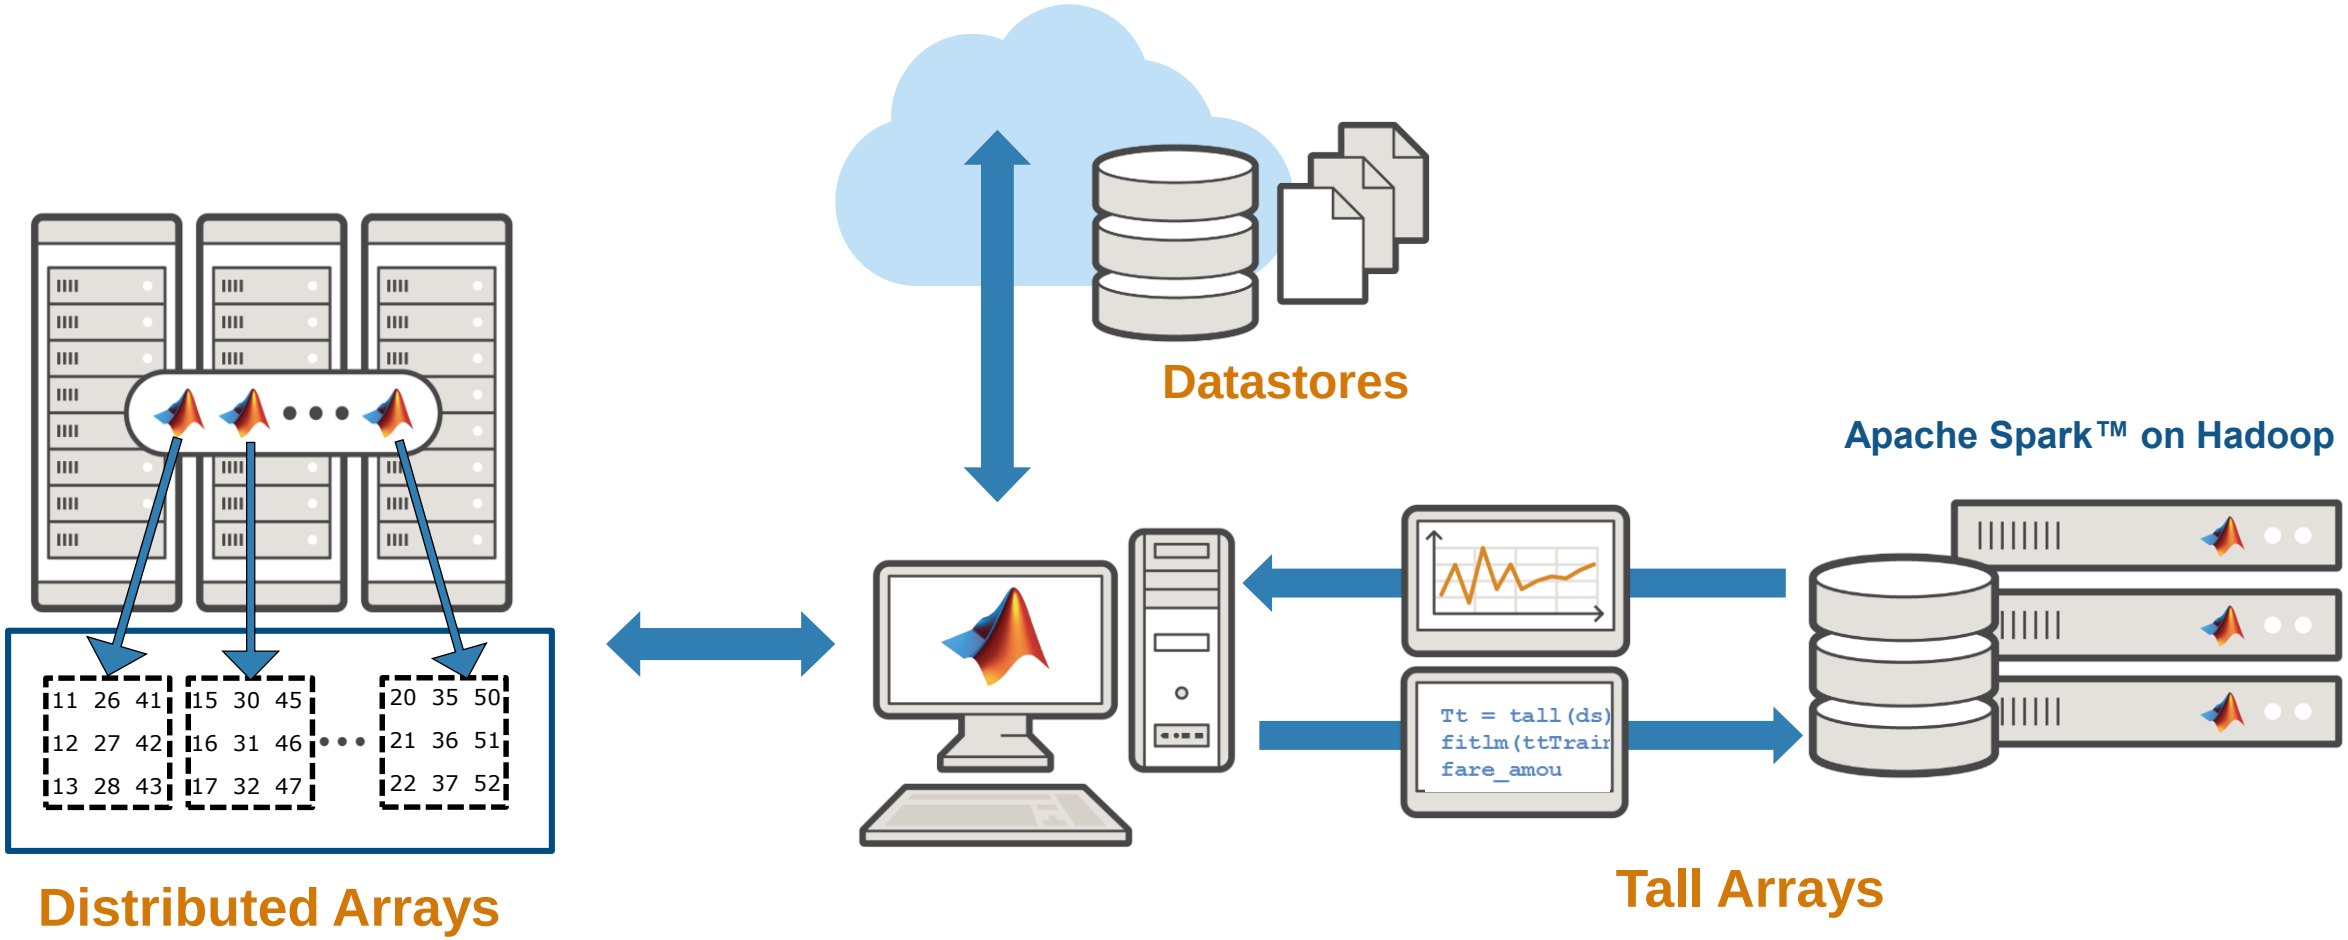

`tall` arrays R2016b

- With Parallel Computing Toolbox, process several “chunks” at once
- Can scale up to clusters with MATLAB Distributed Computing Server

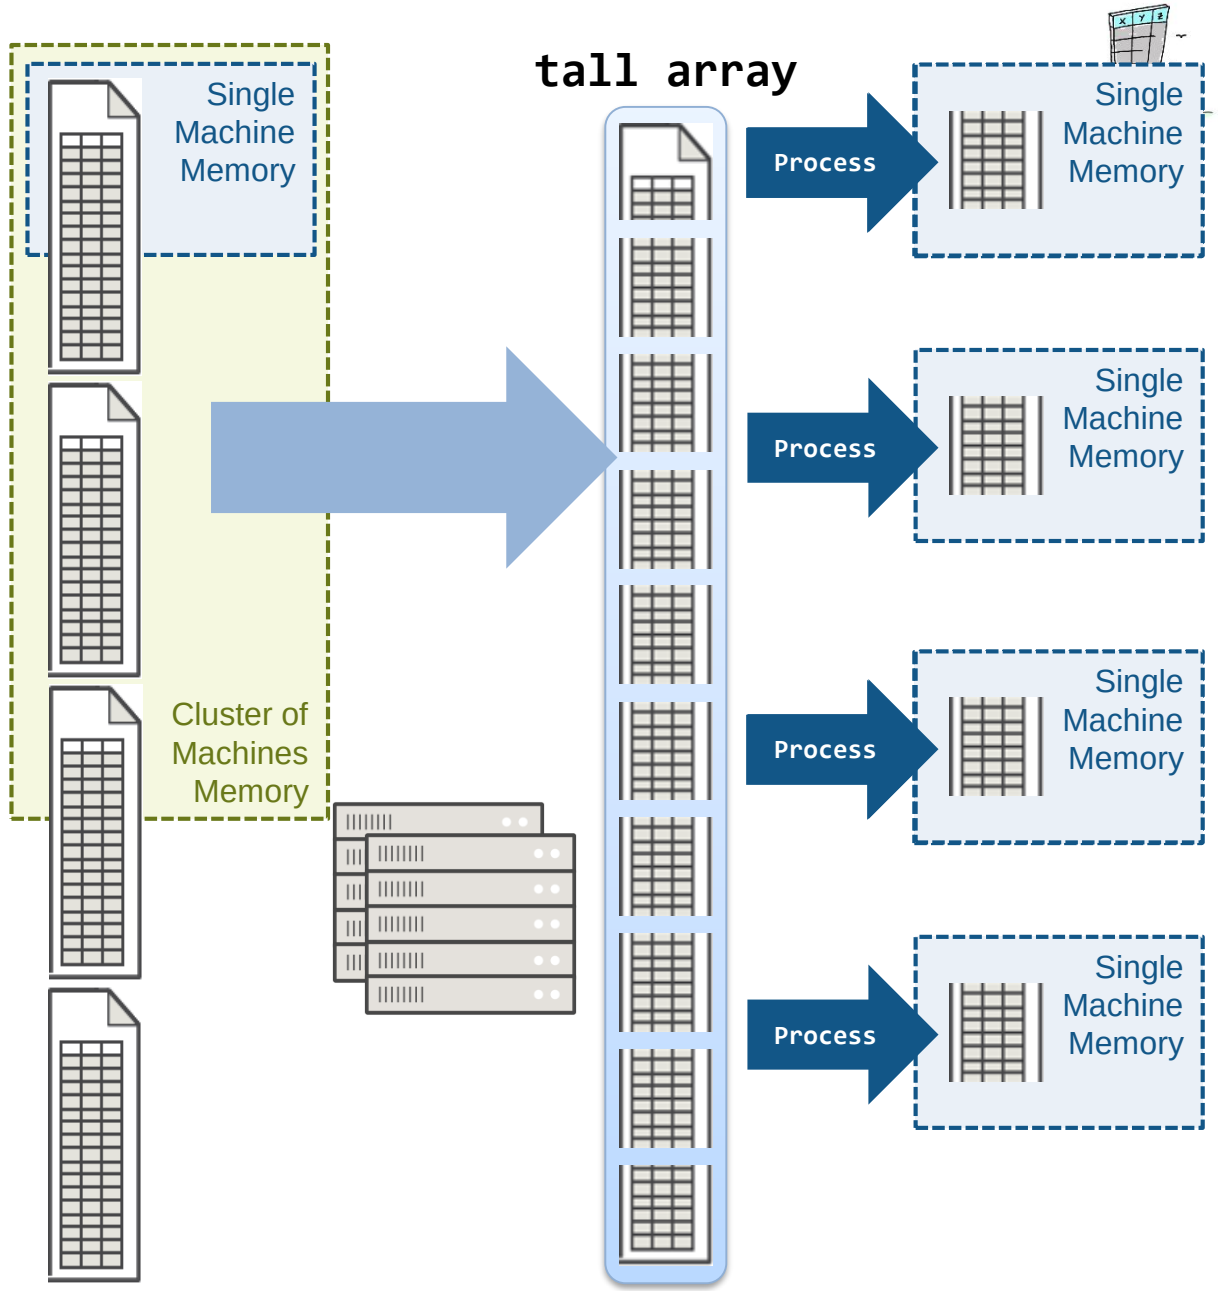

Using Tall Arrays

- **Run in parallel on compute clusters**
  MATLAB Distributed Computing Server

Compute Clusters

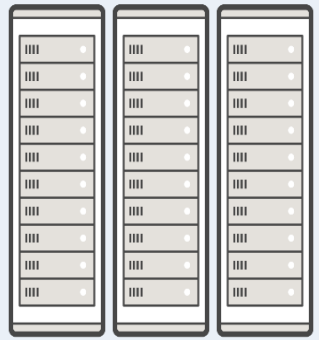

Local disk
Shared folders
Databases

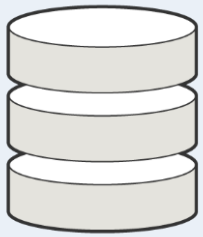

- **Tall arrays**
  MATLAB
- **100's of functions supported**
  MATLAB
  Statistics and Machine Learning Toolbox
- **Run in parallel**
  Parallel Computing Toolbox

Spark + Hadoop

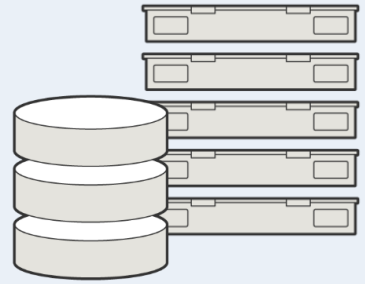

- **Run in parallel on Spark clusters**
  MATLAB Distributed Computing Server
- **Deploy MATLAB applications as standalone applications on Spark clusters**
  MATLAB Compiler

Programming Parallel Applications – Programming Constructs

Parallel-enabled toolboxes

Simple programming constructs
Eg. Parfor, Batch

Advanced programming constructs
Spmd,createJob,labsend

Demo: Getting Data from a Web API

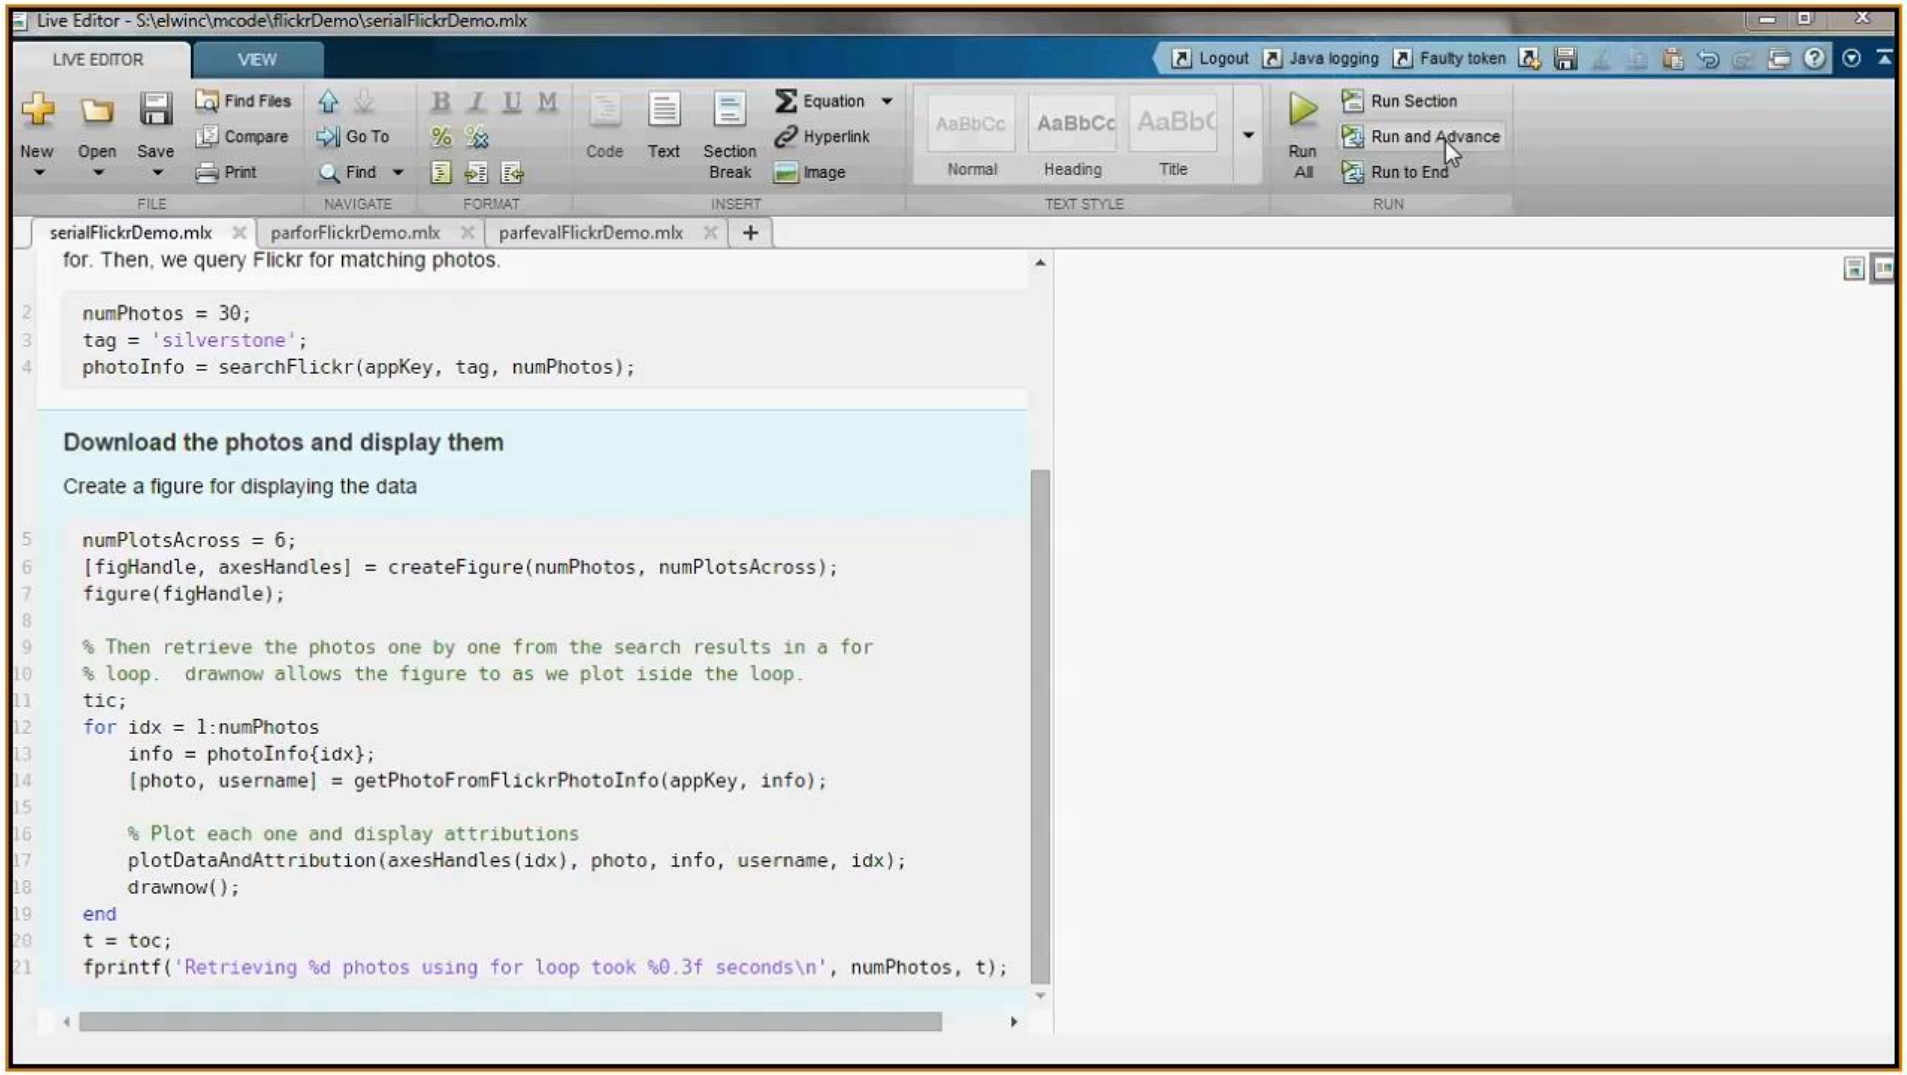

Demo: Getting Data from a Web API using `parfor`

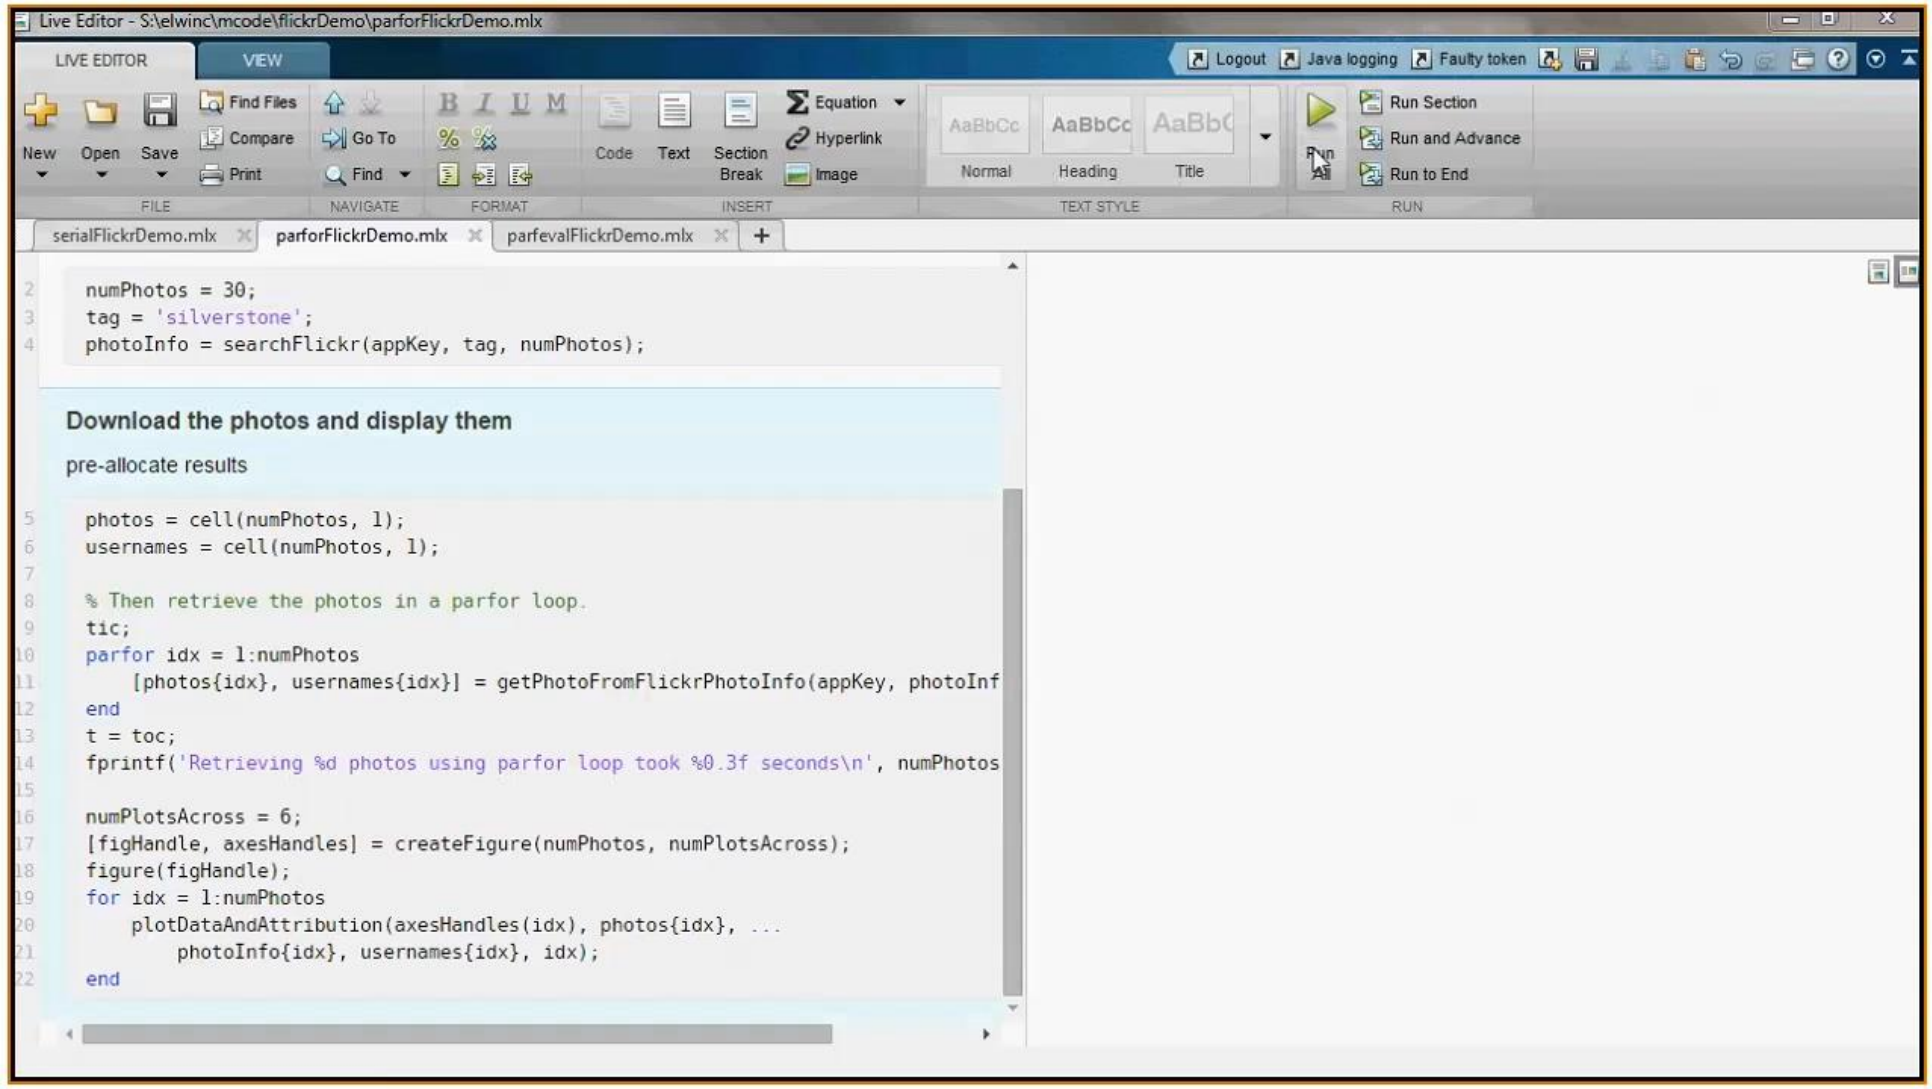

Explicit Parallelism: Independent Tasks or Iterations

Simple programming constructs using `parfor`, `parfeval`

Examples: parameter sweeps, Monte Carlo simulations

No dependencies or communications between tasks

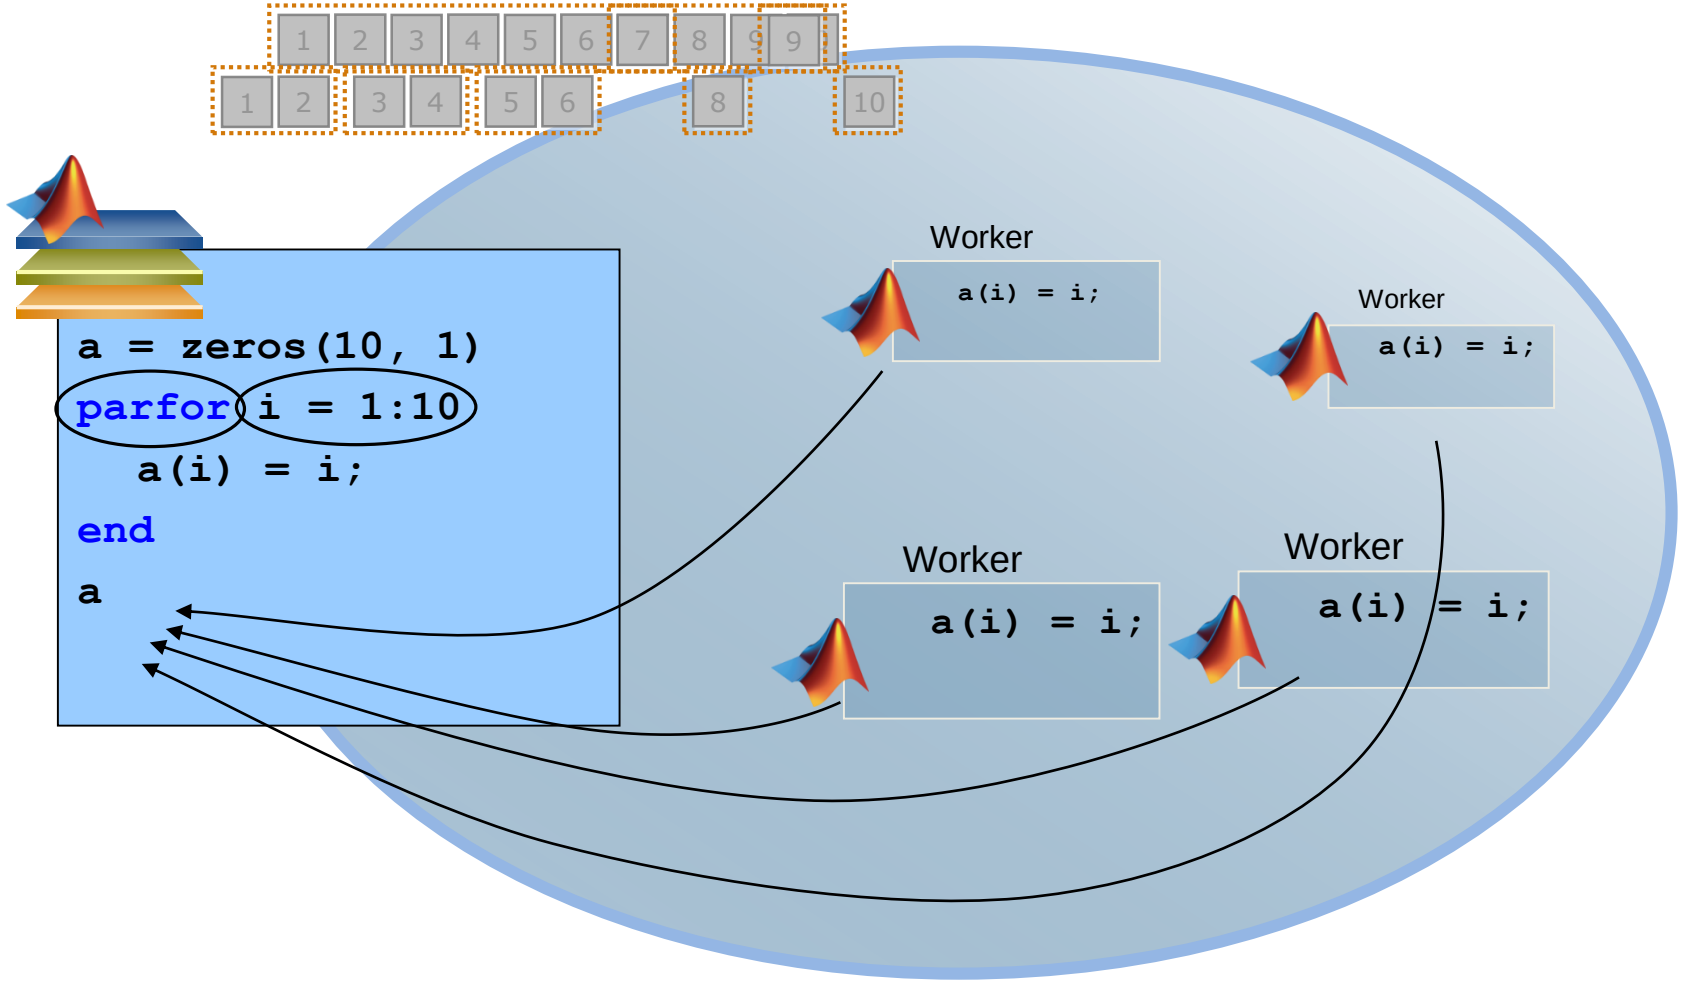

Independent Tasks or Iterations

Simple programming constructs using `parfor`, `parfeval`

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

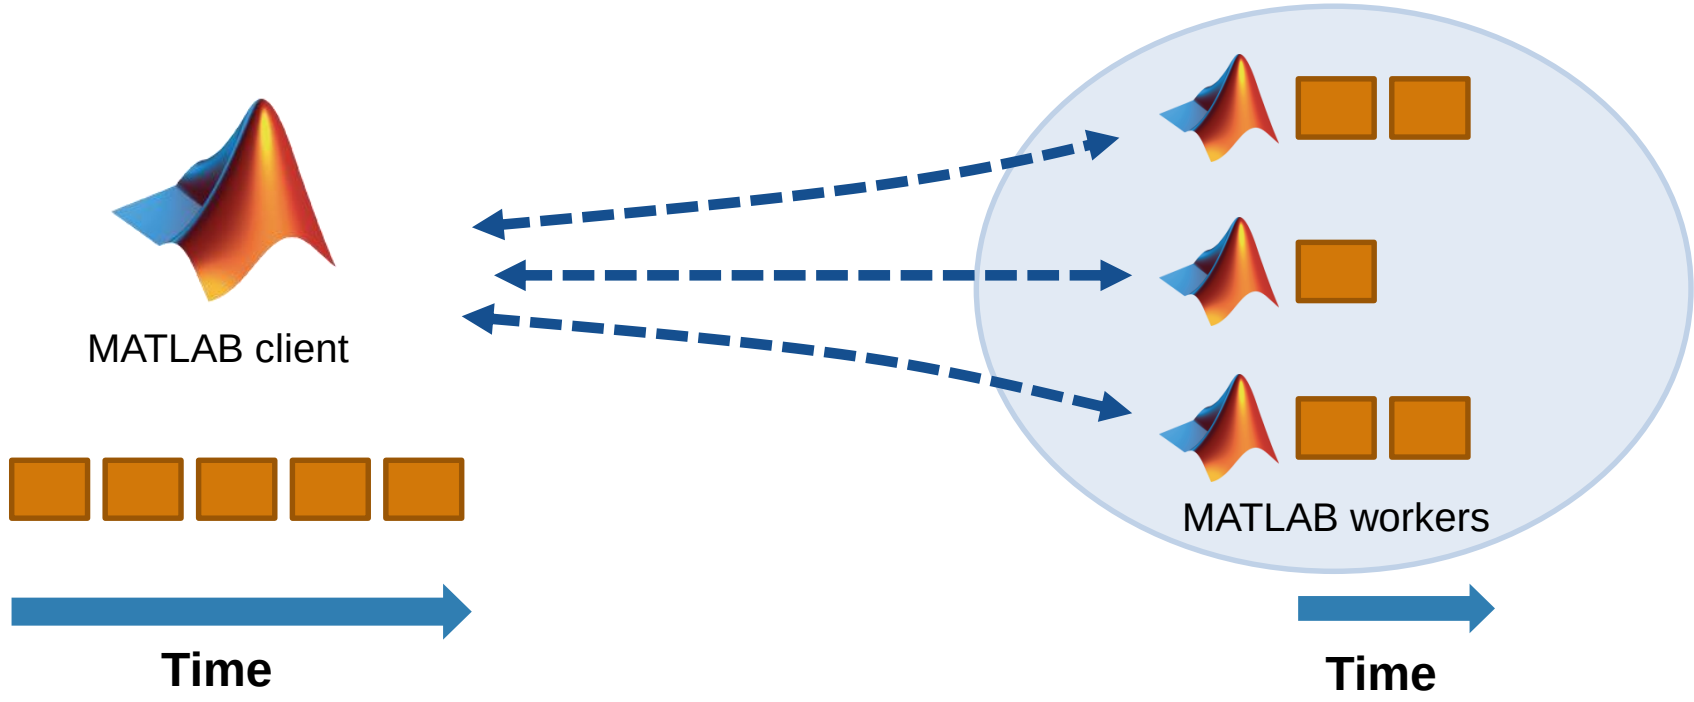

Leverage Parallel Computing for Simulink

- Reduce the total amount of time it takes to…

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

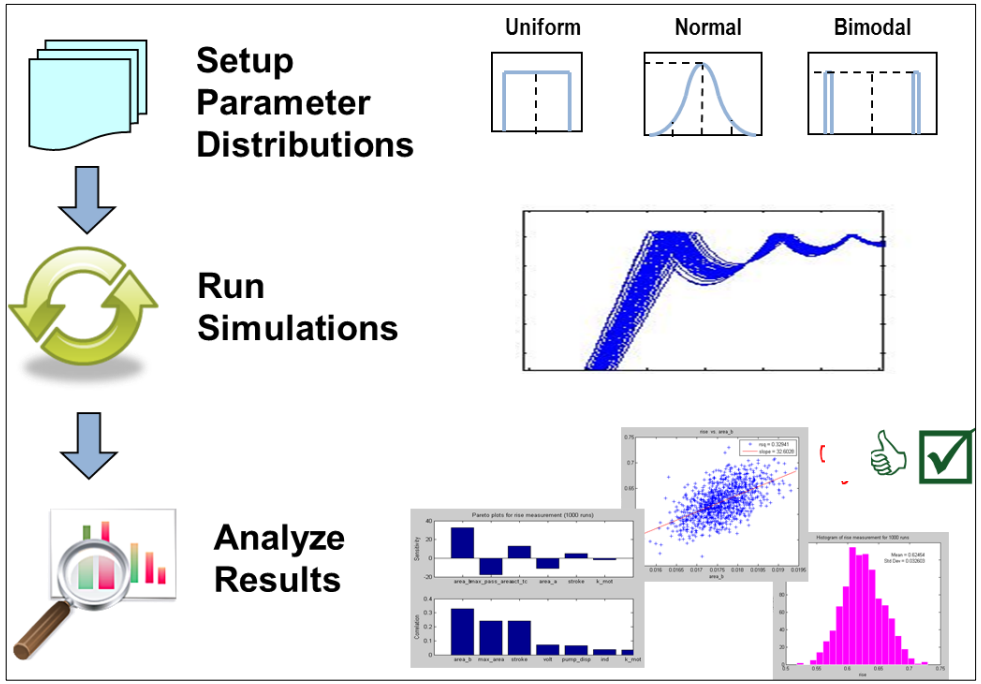

Update models containing large model reference hierarchies

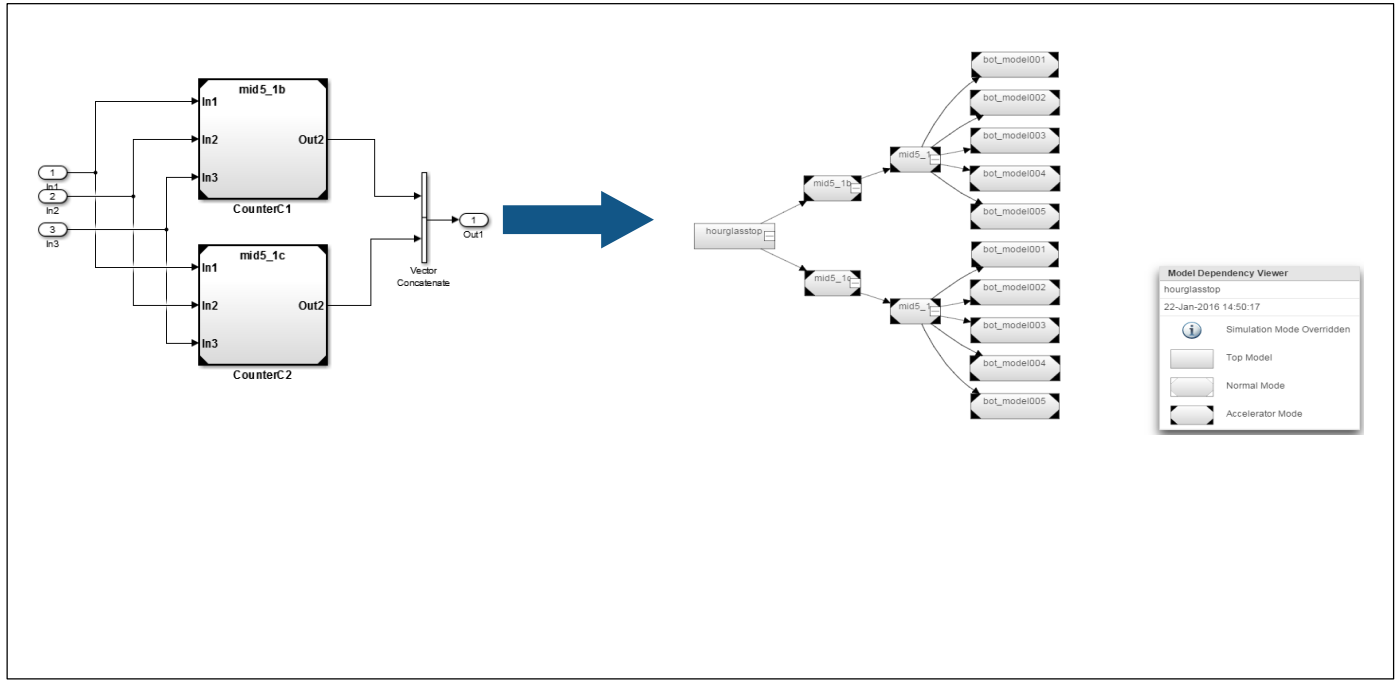

Parallel Simulations using Simulink and Parsim

Directly run multiple parallel simulations from the parsim command

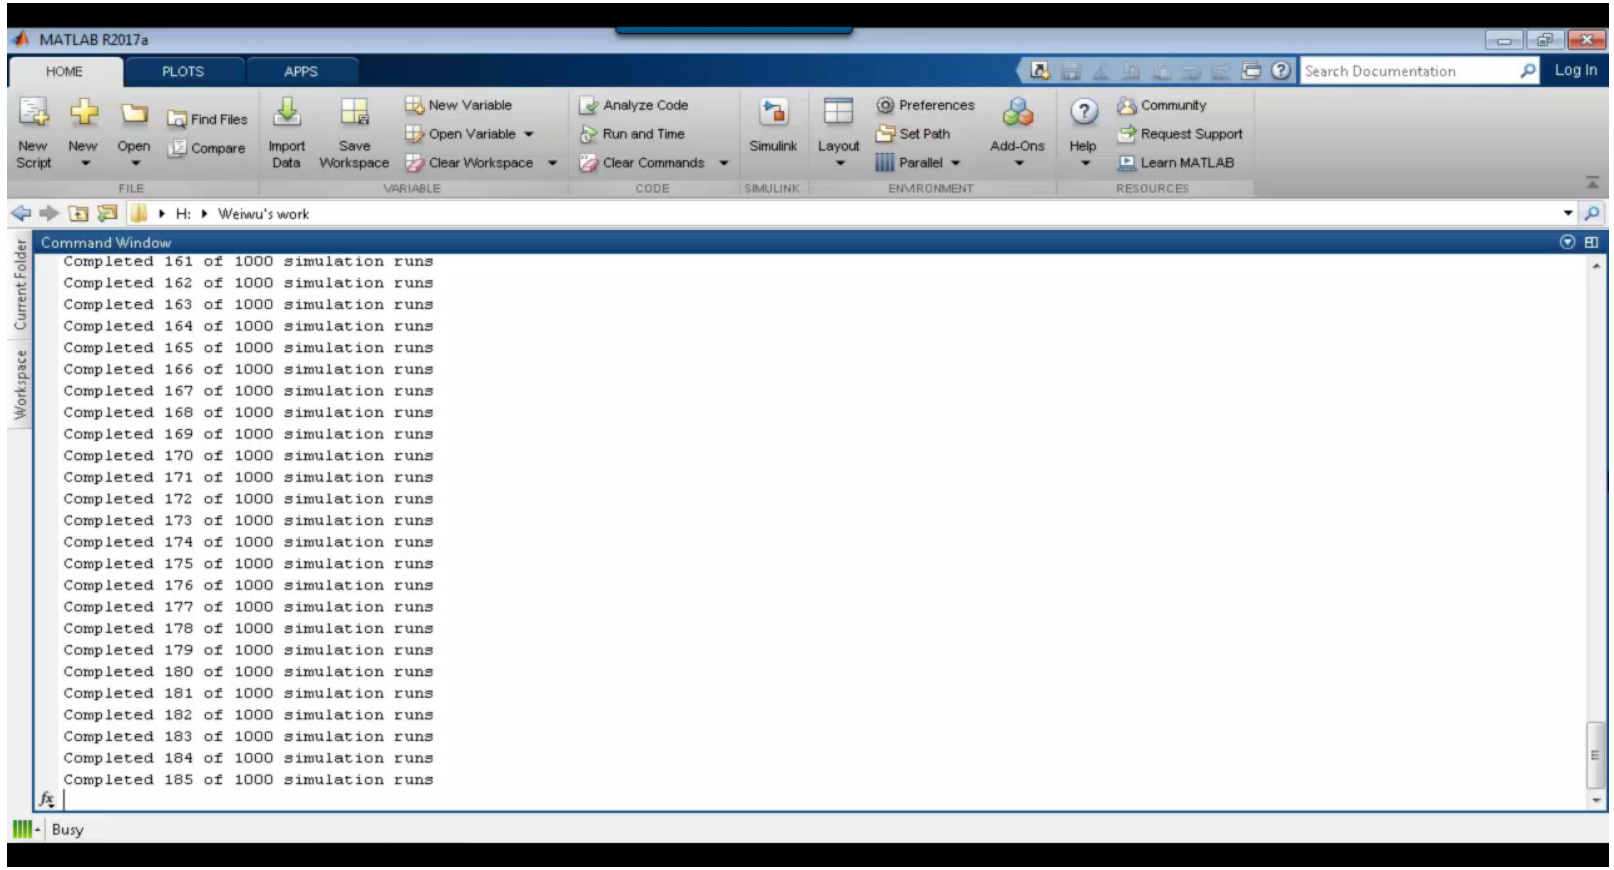

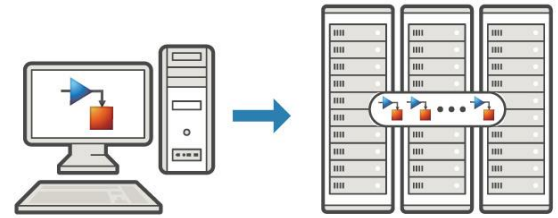

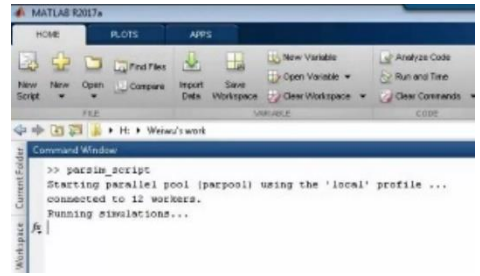

We manage the parallel setup so customers can focus on their simulations

- Enables customers to easily use Simulink with parallel computing
- Simplifies customers' large simulation runs and improves their productivity

Programming Parallel Applications – Programming Constructs

Parallel-enabled toolboxes

Simple programming constructs
Eg. Parfor, Batch

Advanced programming constructs
spmd,createJob,labsend

Agenda

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

Example 3: Solving 2D Wave Equation

GPU Computing

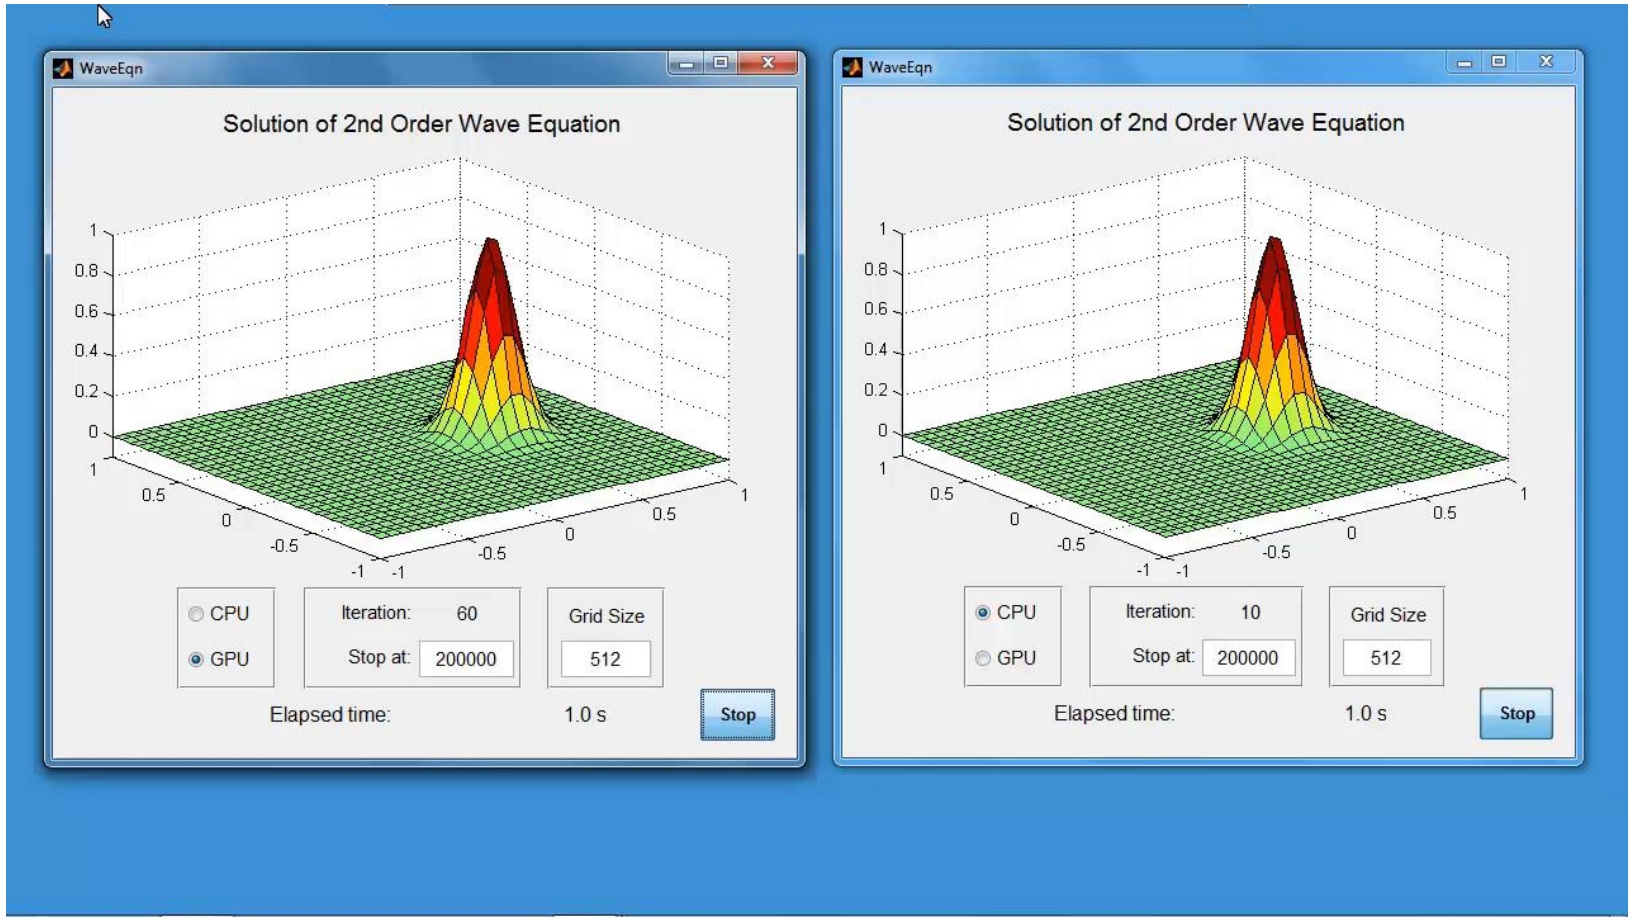

Example: Solving 2D Wave Equation

GPU Computing

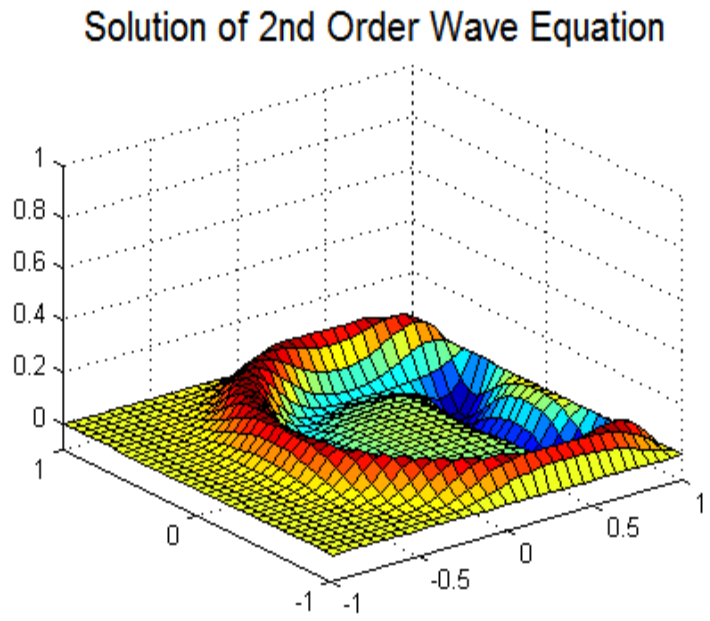

$$\frac{\partial^2 u}{\partial t^2} = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}$$

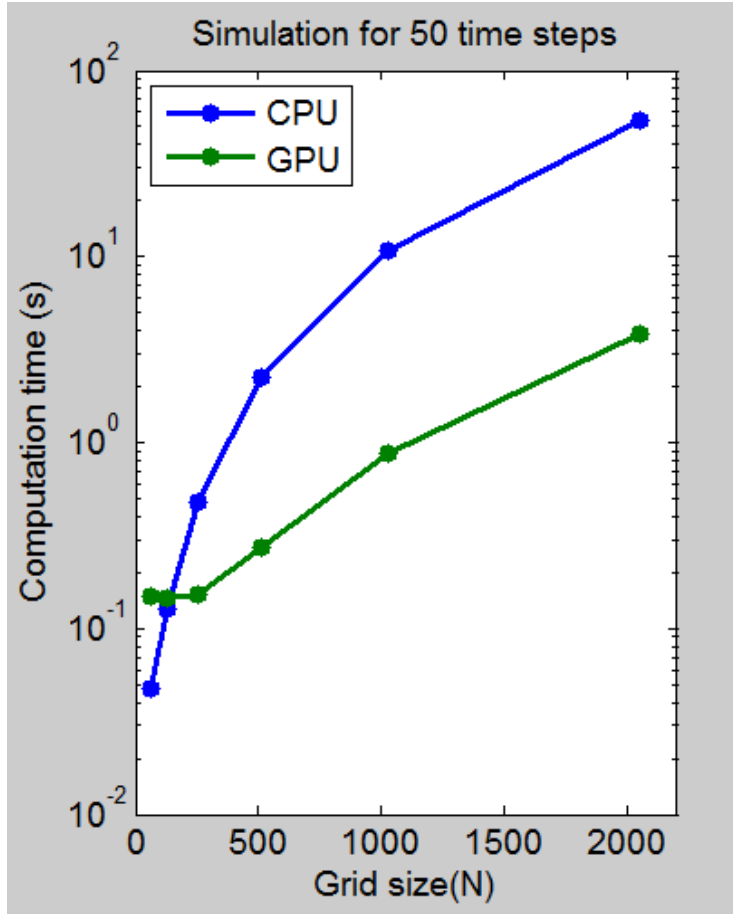

Intel Xeon Processor W3690 (3.47GHz),
NVIDIA Tesla K20 GPU

Parallel Computing Paradigm

Going Parallel: GPUs

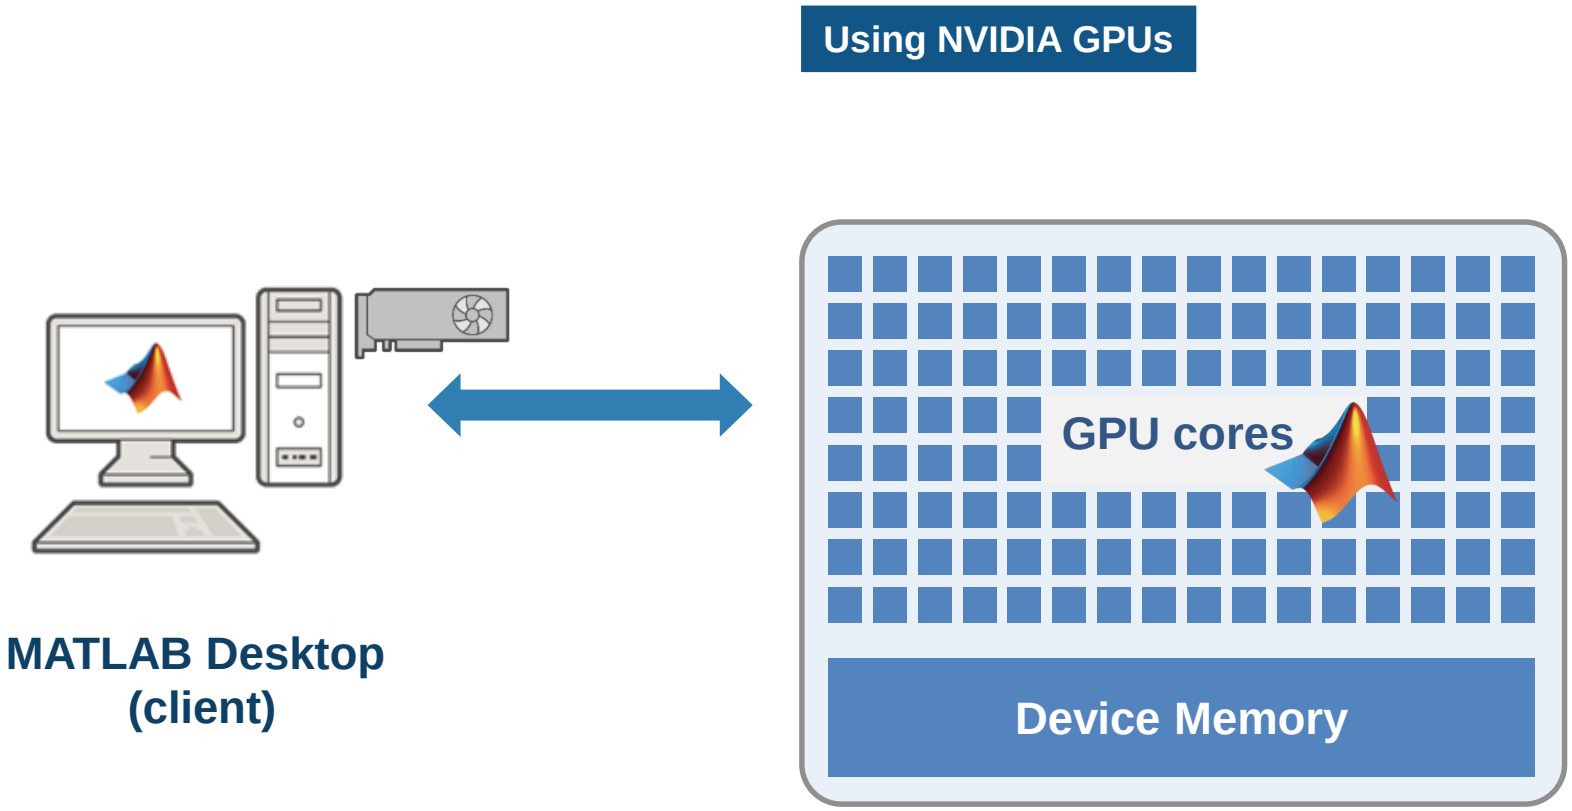

Performance Gain with More Hardware

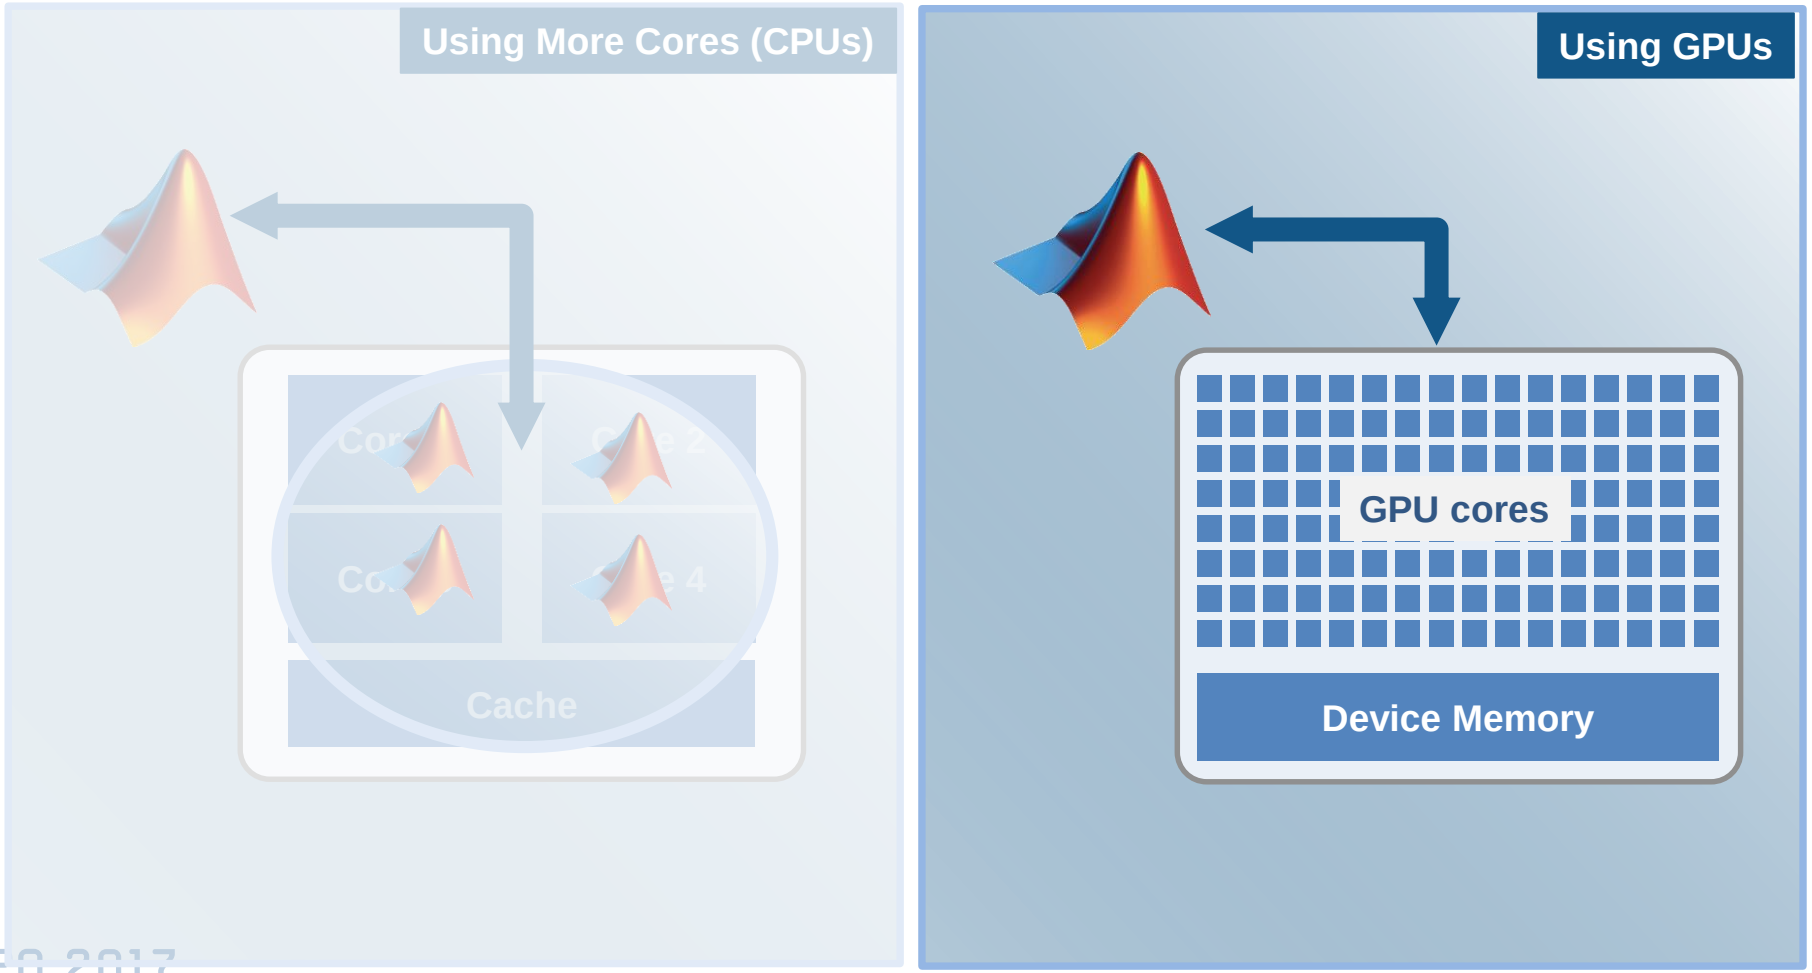

Speed-up using NVIDIA GPUs

- 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

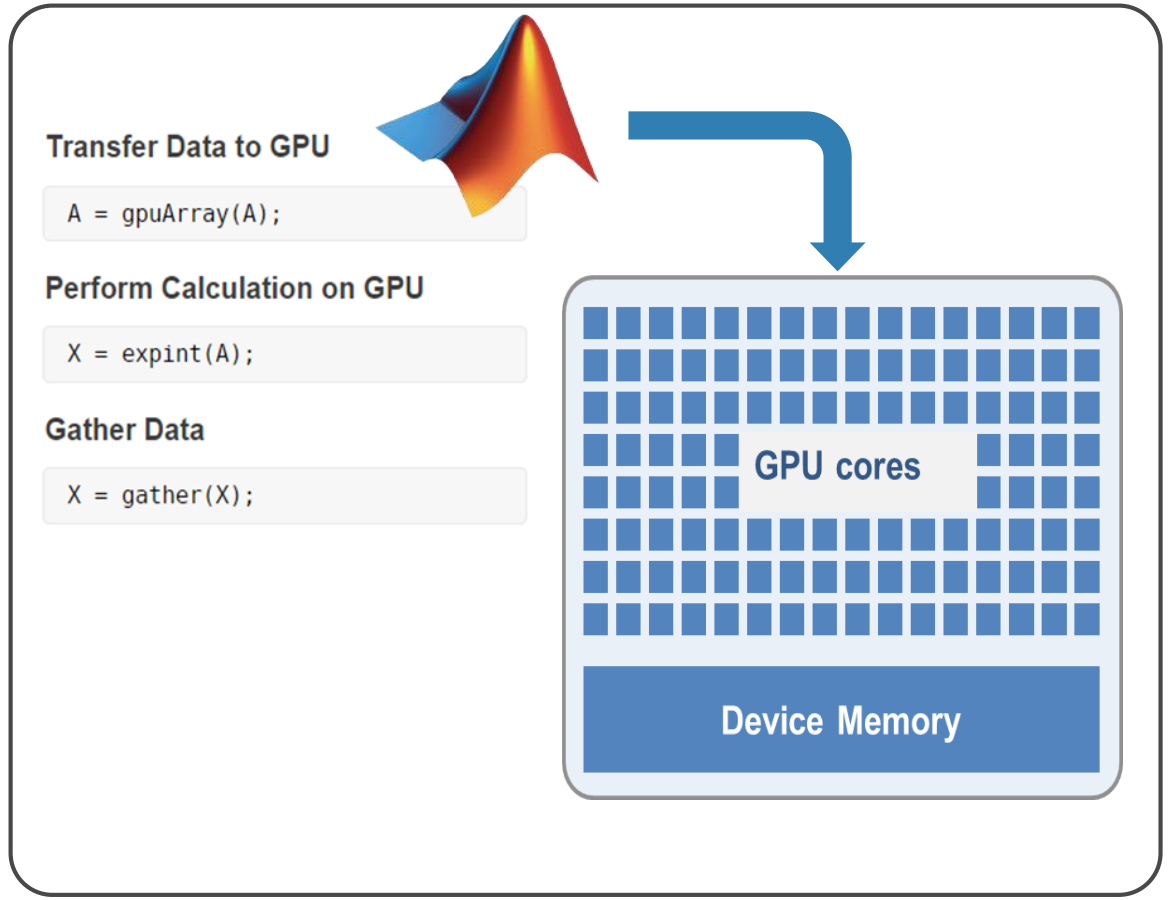

Signal Processing – Acoustic Data Analysis

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 hardware 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

Scaling to a computer cluster

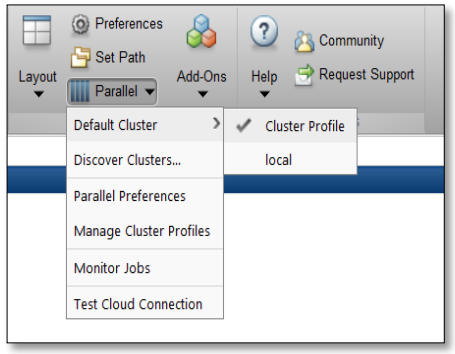

MATLAB Desktop
(client)

Cluster of computers

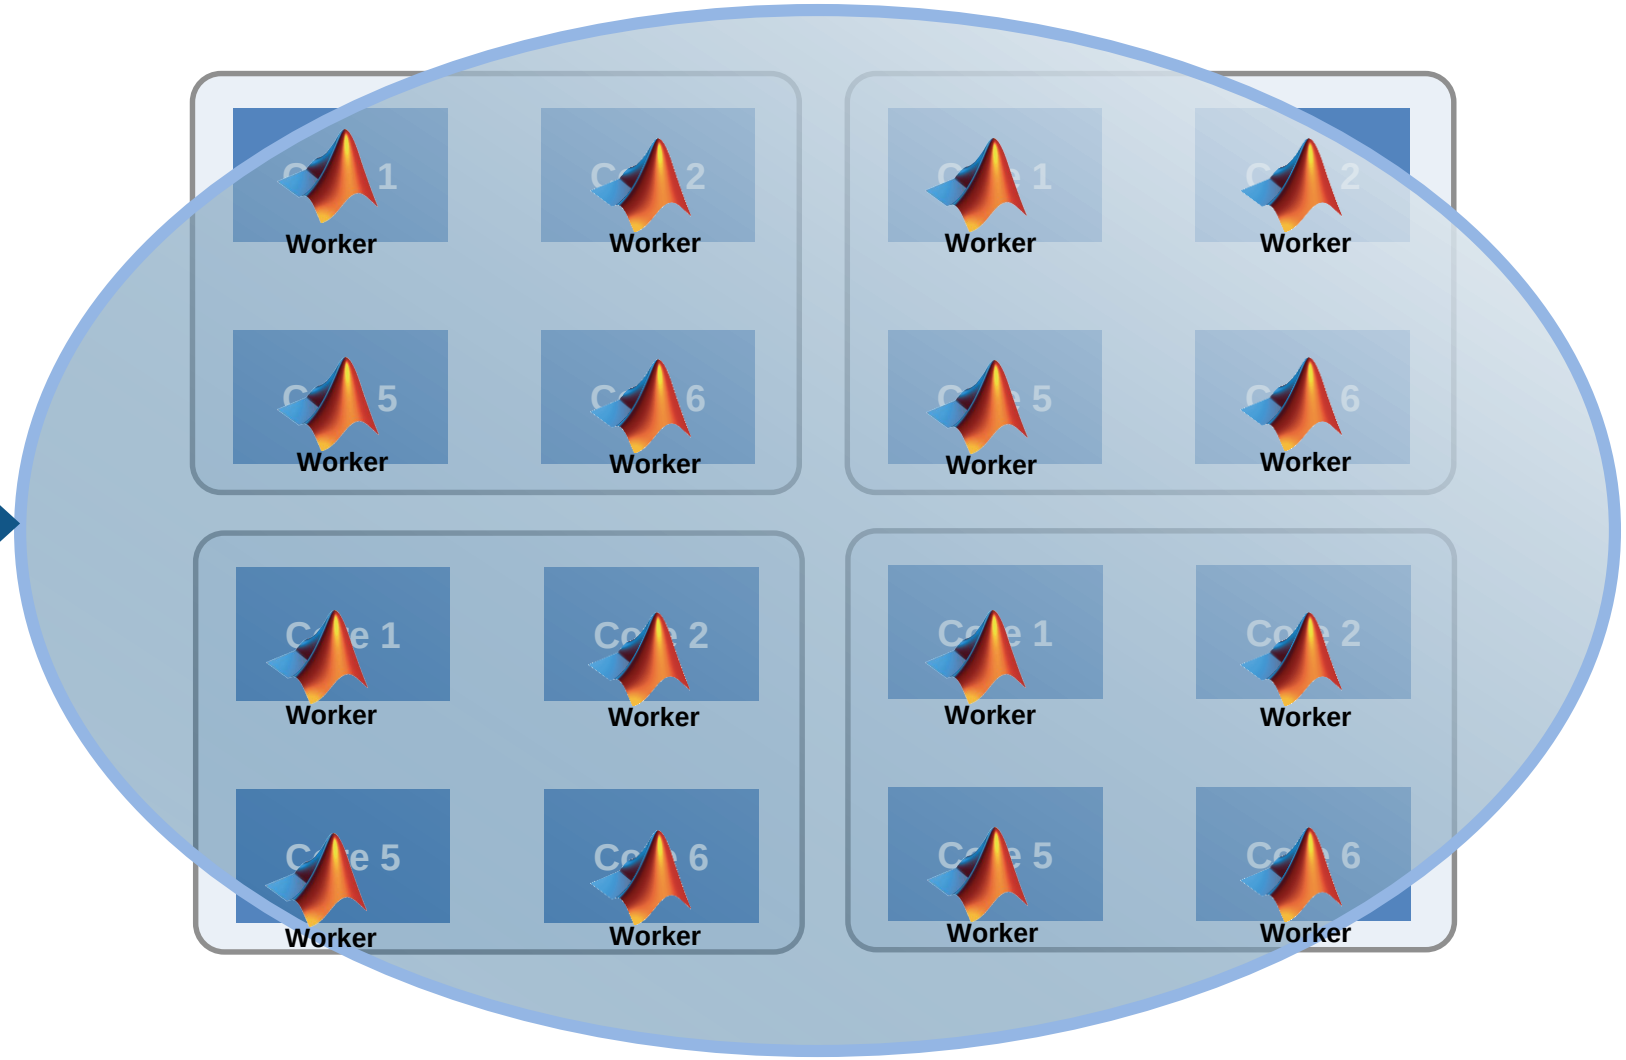

Why parallel computing matters

Scaling case study with a compute cluster

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

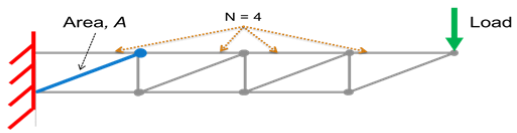

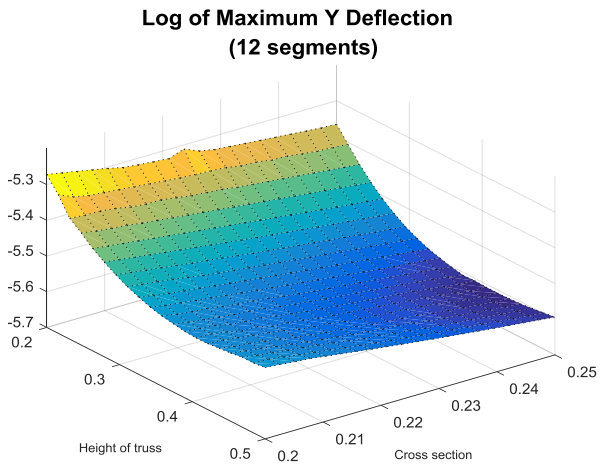

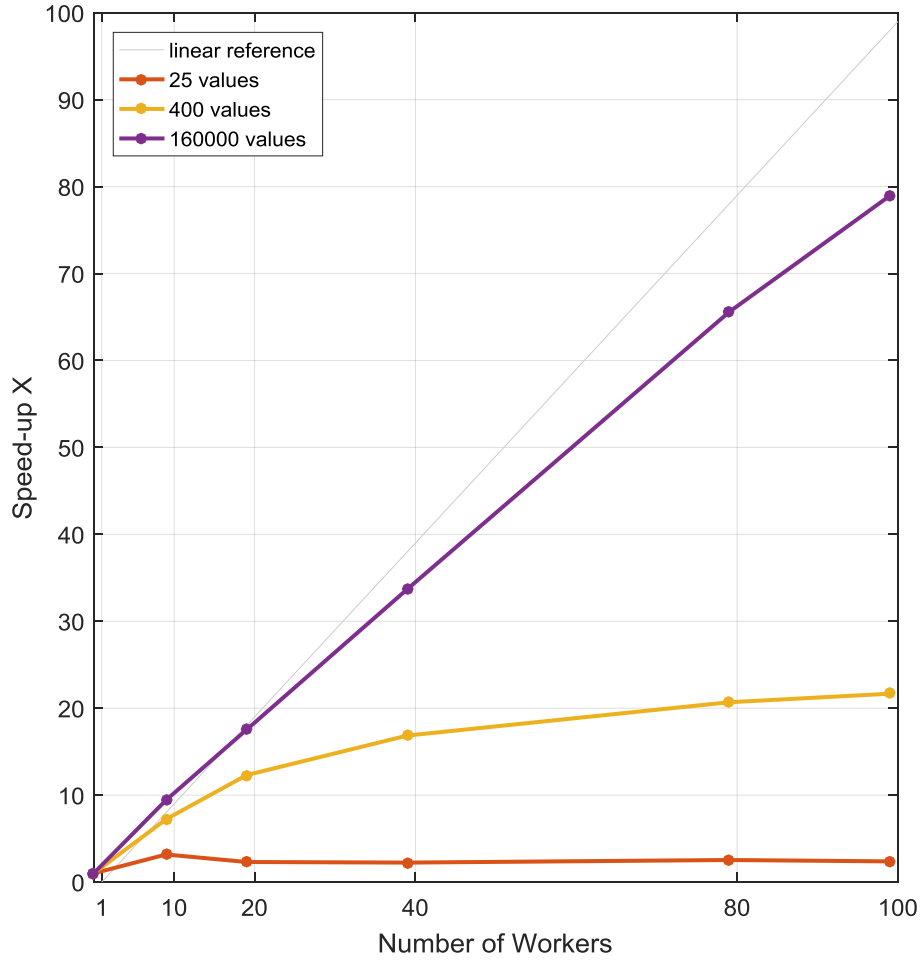

Workers in pool	Compute time (minutes)		
	160e3 values	400 values	25 values
1	140	0.38	0.03
10	15	0.05	0.01
20	8.0	0.03	0.01
40	4.2	0.02	0.01
80	2.1	0.02	0.01
100	1.8	0.02	0.01

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

Scale Up to Clusters, Grids and Clouds

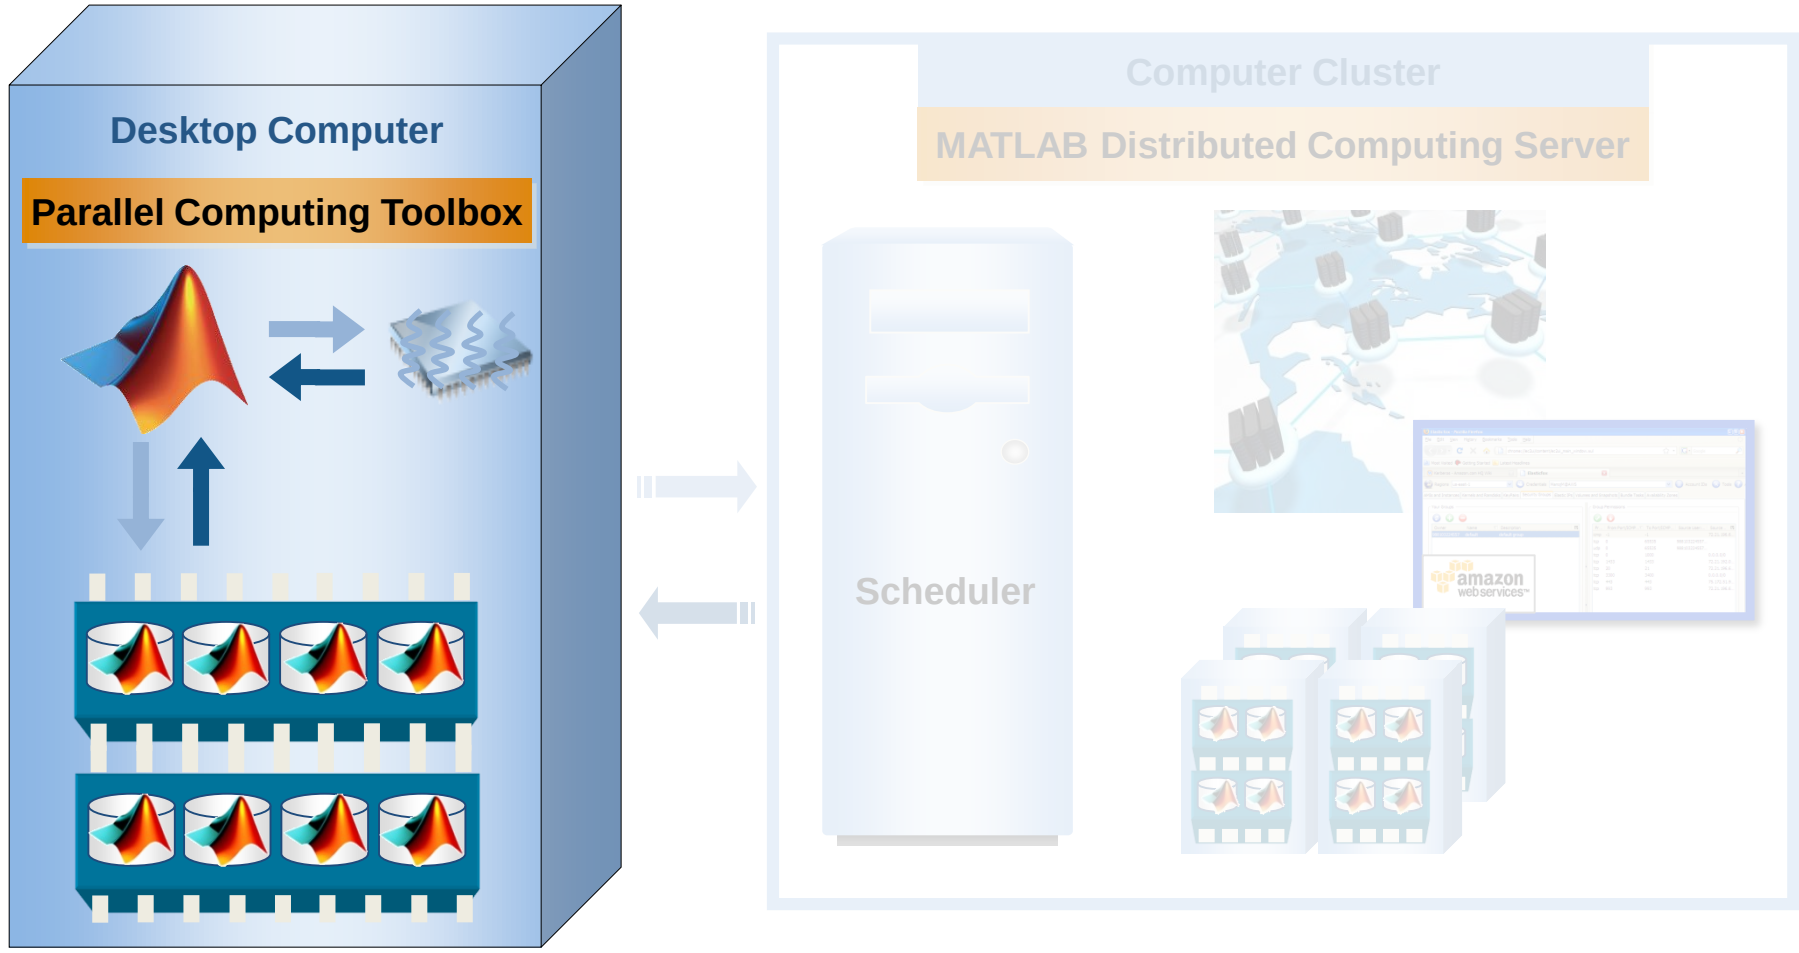

Offload and Scale Computations with `batch`

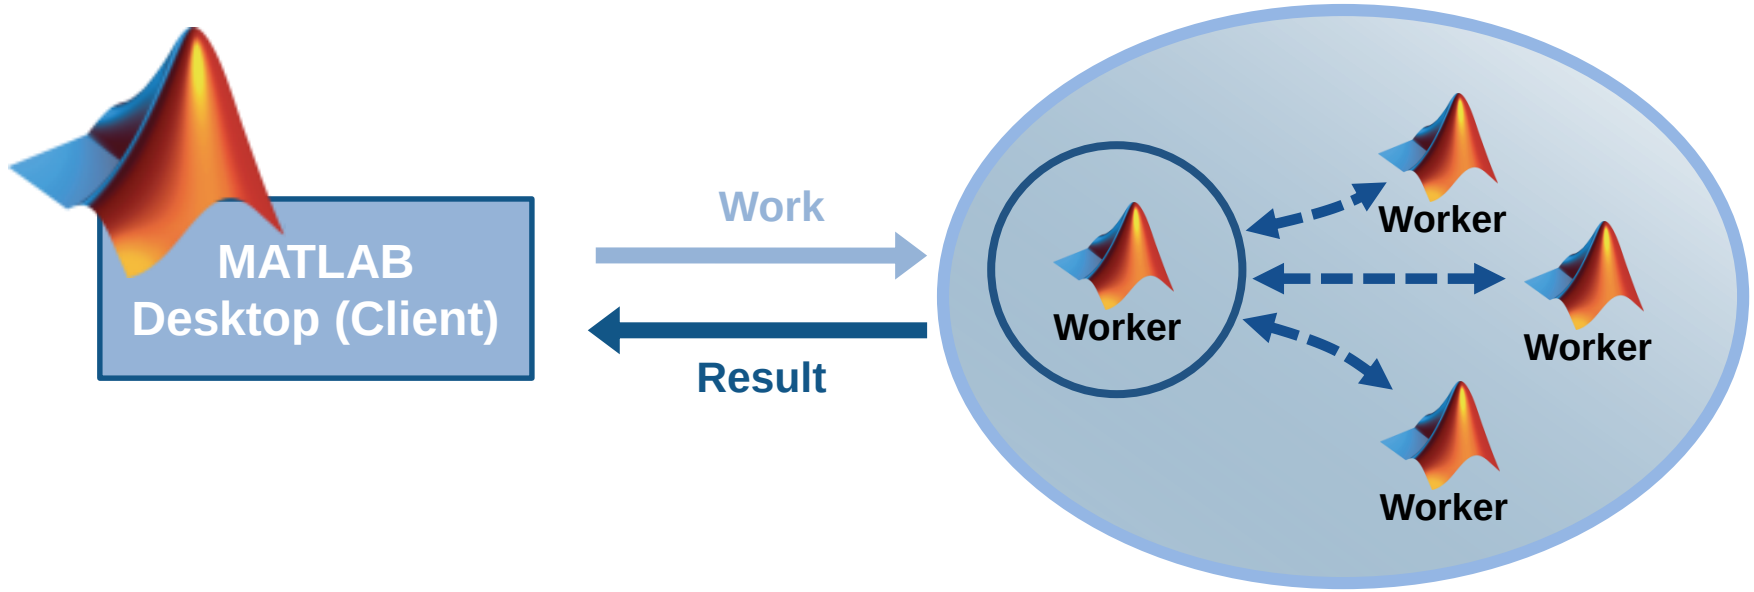


```
batch(…, 'Pool',…)
```

Scale your applications beyond the desktop

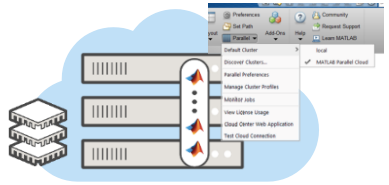

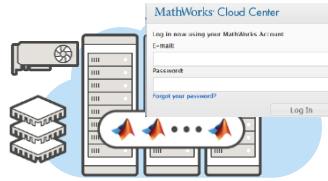

Option	Parallel Computing Toolbox	MATLAB Parallel Cloud	MATLAB Distributed Computing Server *for Amazon EC2*	MATLAB Distributed Computing Server *for Custom Cloud*	MATLAB Distributed Computing Server
Description	Explicit desktop scaling	Single-user, basic scaling to cloud	Scale to EC2 with some customization	Scale to custom cloud	Scale to clusters
Maximum workers	No limit	16	256	No limit	No limit
Hardware	Desktop	MathWorks Compute Cloud	Amazon EC2	Amazon EC2, Microsoft Azure, Others	Any
Availability	Worldwide	United States and Canada	United States, Canada and other select countries in Europe	Worldwide	Worldwide

Learn More: Parallel Computing on the Cloud

Summary and Takeaways

- **Speed up your MATLAB and Simulink applications without being an expert**
  - Reduce computation time by using more cores or accessing Graphical Processing Units
- **Leverage the Parallel Computing Toolbox to**
  - Reduce Computation Time: `parfor`, `gpuArray`, `parsim`
  - Offload and Scale Computations: `batch`
- **Speed up Big Data Analytics**
  - Using `datastore` and `tall` or `distributed` arrays
- **Develop parallel applications on the desktop and scale to clusters seamlessly**

What's new in 16b and 17a?

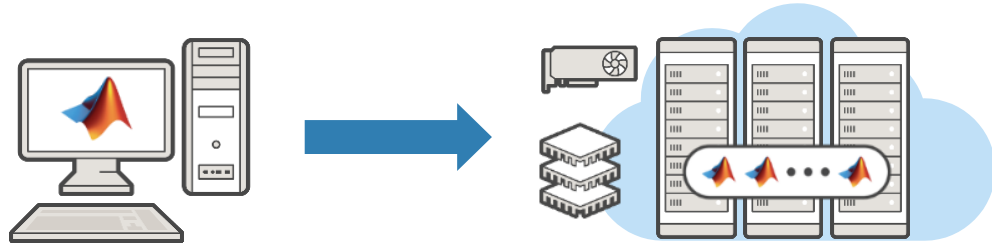

R2016b

- `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 Training Offerings

MATLAB Programming Techniques

INTERMEDIATE

This two-day course covers details of performance optimization as well as tools for writing, debugging, and profiling code. Topics include:

- Creating robust applications
- Structuring code
- Structuring data
- Creating custom toolboxes

Prerequisites: *MATLAB Fundamentals*

Parallel Computing with MATLAB

INTERMEDIATE

This two-day course shows how to use Parallel Computing Toolbox™ to speed up existing code and scale up across multiple computers using MATLAB Distributed Computing Server™ (MDCS). Attendees who are working with long-running simulations, or large data sets, will benefit from the hands-on demonstrations and exercises in the course. Topics include:

- Parallel for-loops
- Offloading execution
- Working with clusters
- Distributing and processing large data sets
- GPU computing

Prerequisites: *MATLAB Fundamentals*

http://www.mathworks.com/services/training/

Learn Further

- https://www.mathworks.com/solutions/parallel-computing.html
- https://www.mathworks.com/help/distcomp/
- https://www.mathworks.com/help/distcomp/examples.html

Speaker Details

Email: Alka.Nair@mathworks.in

Contact MathWorks India

Products/Training Enquiry Booth

Call: 080-6632-6000

Email: info@mathworks.in

Your feedback is valued.

Please complete the feedback form provided to you.


MathWorks®