MATLAB EXPO 2016
KOREA

4월 28일 (목)

등록 하기 matlabexpo.co.kr

- Data analytics Workflow
- Map reduce demo refine. (Car register or weather station)

Working with Big Data using MATLAB

성호현 차장
Senior Application Engineer
MathWorks Korea

Challenges of Data

"Any collection of data sets so large and complex that it becomes difficult to process using … traditional data processing applications."
(Wikipedia)

- Various Data Sources
- Rapid data exploration
- Development of scalable algorithms
- Ease of deployment

Agenda

- How big is big?
- Reading big data
- Processing quite big data
- Processing big data
- Summary

How big is big?

What does “Big Data” even mean?

“Any collection of data sets so large and complex that it becomes difficult to process using … traditional data processing applications.”

(Wikipedia)

“Any collection of data sets so large that it becomes difficult to process using traditional MATLAB functions, which assume all of the data is in memory.”

(MATLAB)

How big is big?

Not a new problem

- In 1085 William 1st commissioned a survey of England
  - ~2 million words and figures collected over two years
  - too big to handle in one piece
  - collected and summarized in several pieces
  - used to generate revenue (tax), but most of the data then sat unused

How big is big?

A new problem

- The Large Hadron Collider was switched back on earlier this year
  - ~600 million collisions per second (only a fraction get recorded)
  - amounts to 30 petabytes per year
  - too big to even store in one place
  - used to explore interesting science, but taking researchers a long time to get through

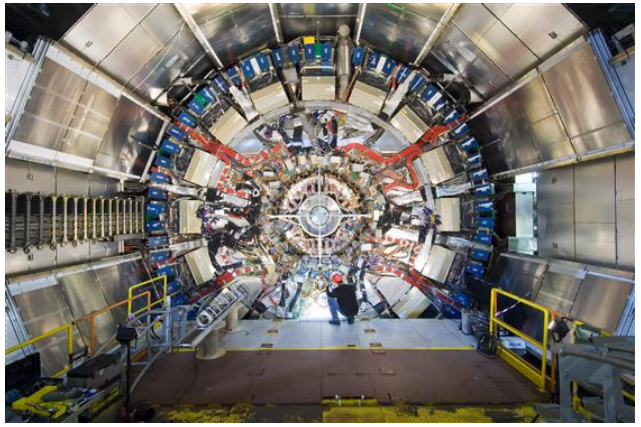

Image courtesy of CERN. Copyright 2011 CERN.

How big is big?

Sizes of data in this talk

- Most of our data lies somewhere in between
  - a few MB up to a few TB
  - <1GB can typically be handled in memory on one machine (small data)
  - 1-100GB can typically be handled in memory of many machines (quite big data)
  - >100GB typically requires processing in pieces using many machines (big data)

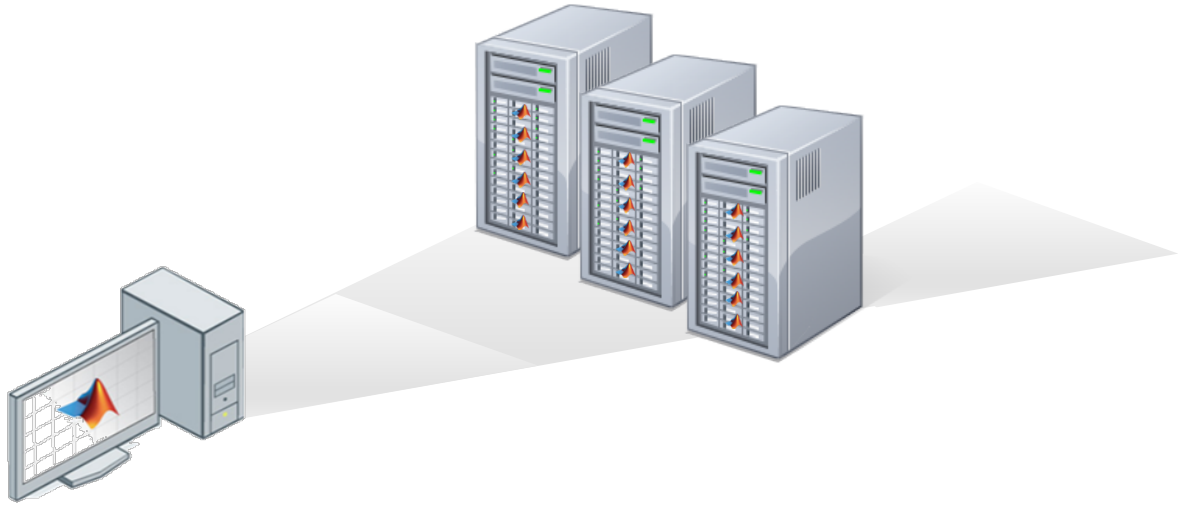

Agenda

- How big is big?
- Reading big data
- Processing quite big data
- Processing big data
- Summary

Reading big data

What tools are there?

Reading big data

What tools are there?

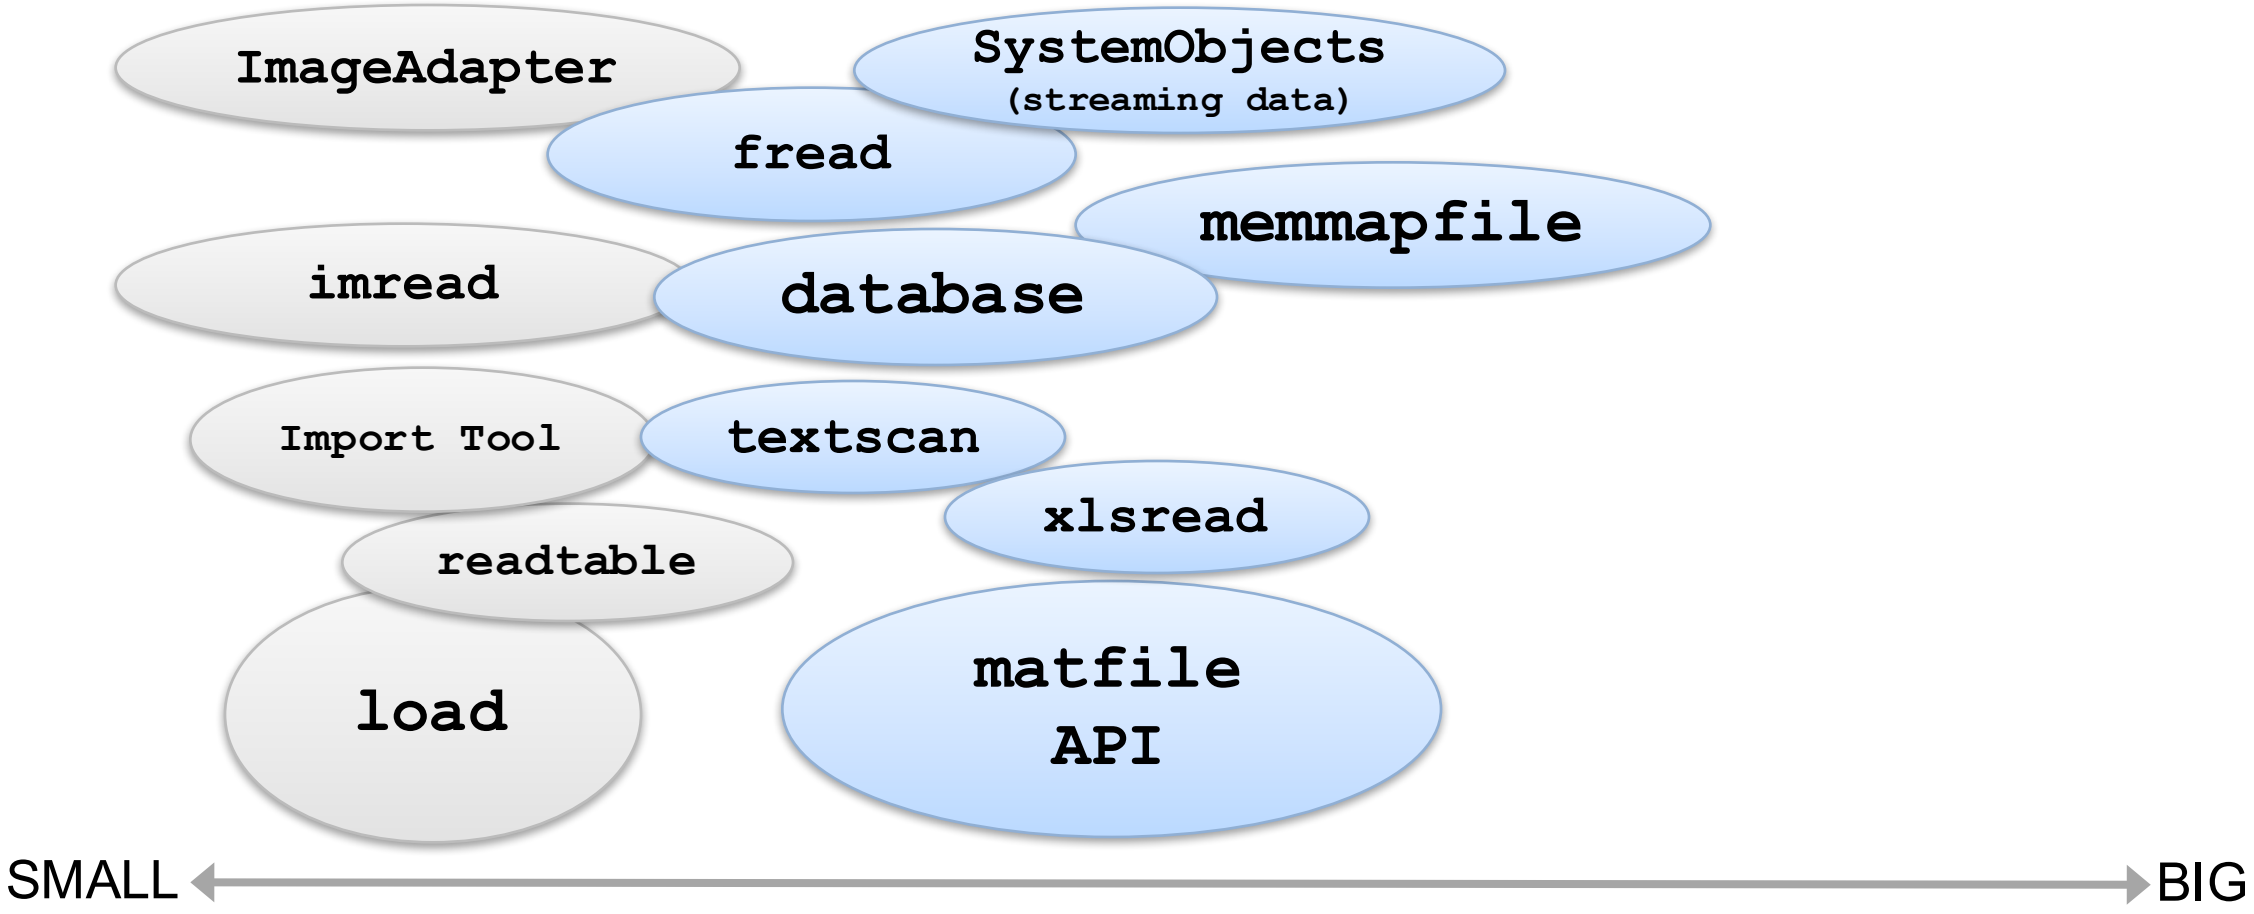

Reading big data

What tools are there?

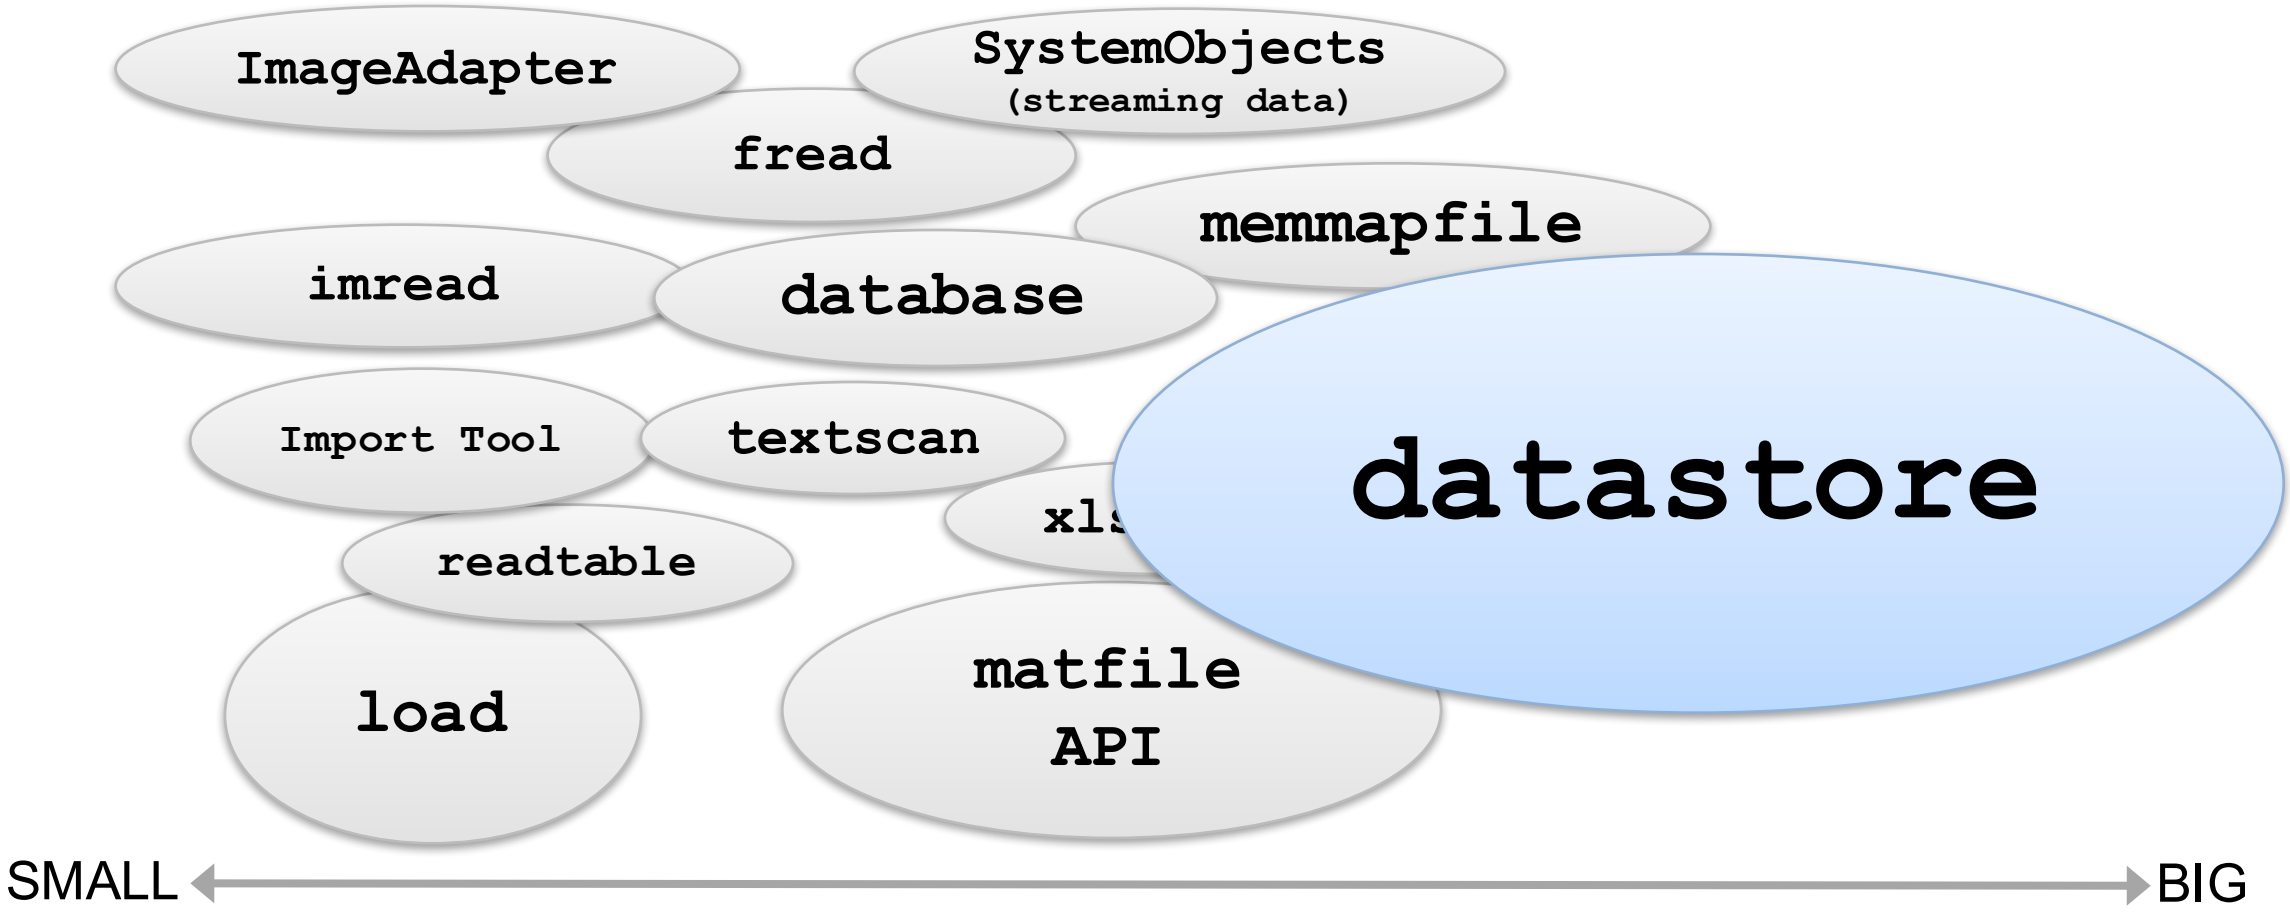

Reading big data

Datastore:

- Simple interface for data in multiple files/folders
- Presents data a piece at a time
- Access pieces in serial (desktop) or in parallel (cluster)
- Back-ends for tabular text, images, databases and more

Reading big data

Datastore DEMO

Agenda

- How big is big?
- Reading big data
- Processing quite big data
- Processing big data
- Summary

Processing quite big data

When the data fits in cluster memory

- Using distributed arrays
  - Use the memory of multiple machines as though it was your own
  - Client sees a “normal” MATLAB variable
  - Work happens on cluster

Processing quite big data

Distributed array functions

- Many common MATLAB functions supported: (about 250)
- Includes most linear algebra
- Scale up your maths

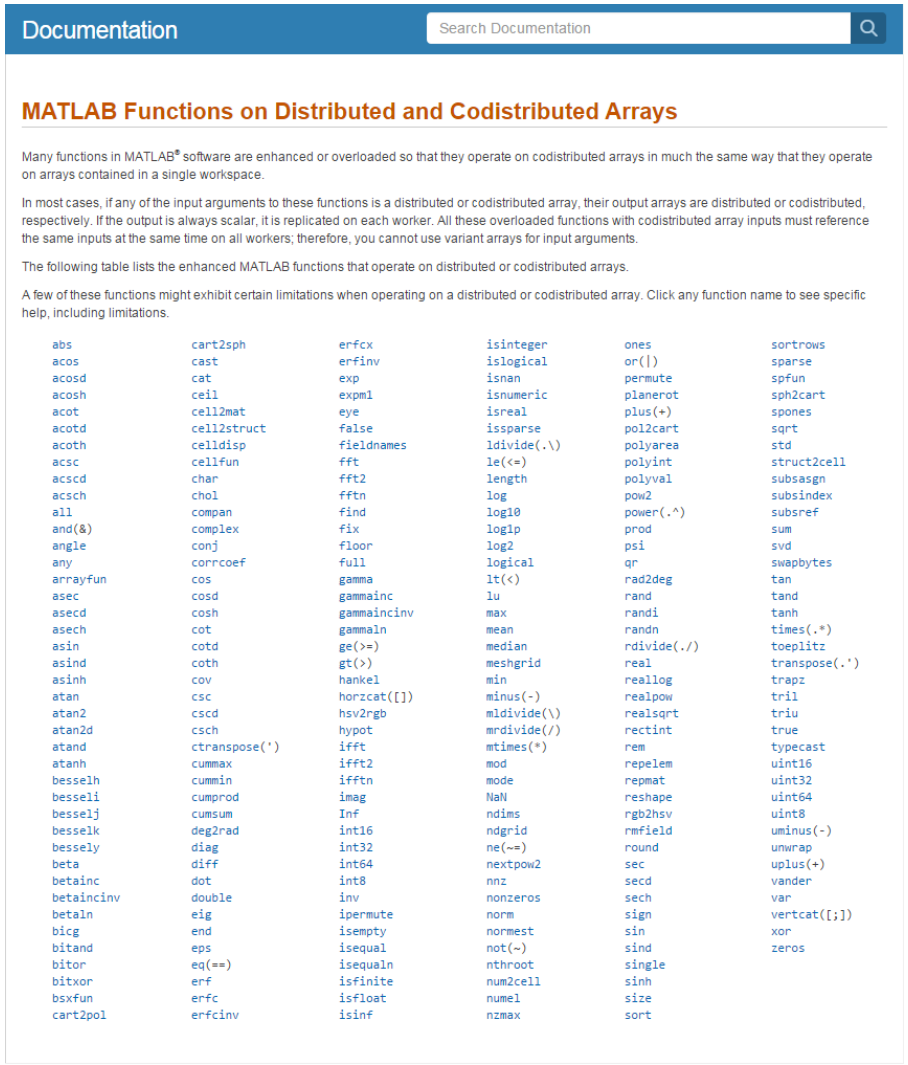

Documentation

Search Documentation

MATLAB Functions on Distributed and Codistributed Arrays

Many functions in MATLAB® software are enhanced or overloaded so that they operate on codistributed arrays in much the same way that they operate on arrays contained in a single workspace.

In most cases, if any of the input arguments to these functions is a distributed or codistributed array, their output arrays are distributed or codistributed, respectively. If the output is always scalar, it is replicated on each worker. All these overloaded functions with codistributed array inputs must reference the same inputs at the same time on all workers; therefore, you cannot use variant arrays for input arguments.

The following table lists the enhanced MATLAB functions that operate on distributed or codistributed arrays.

A few of these functions might exhibit certain limitations when operating on a distributed or codistributed array. Click any function name to see specific help, including limitations.

abs	cart2sph	erfcx	isinteger	ones	sortrows
acos	cast	erfinv	islogical	or(|)	sparse
acosd	cat	exp	isnan	permute	spfun
acosh	ceil	expm1	isnumeric	planerot	sph2cart
acot	cell2mat	eye	isreal	plus(+)	spones
acotd	cell2struct	false	issparse	pol2cart	sqrt
acoth	celldisp	fieldnames	ldivide(.\)	polyarea	std
acsc	cellfun	fft	le(<=)	polyint	struct2cell
acscd	char	fft2	length	polyval	subsasgn
acsch	chol	fftn	log	pow2	subsindex
all	compan	find	log10	power(.^)	subsref
and(&)	complex	fix	log1p	prod	sum
angle	conj	floor	log2	psi	svd
any	corrcoef	full	logical	qr	swapbytes
arrayfun	cos	gamma	lt(<)	rad2deg	tan
asec	cosd	gammainc	lu	rand	tand
asecd	cosh	gammaincinv	max	randi	tanh
asech	cot	gammaln	mean	randn	times(.*)
asin	cotd	ge(>=)	median	rdivide(./)	toeplitz
asind	coth	gt(>)	meshgrid	real	transpose(.')
asinh	cov	hankel	min	reallog	trapz
atan	csc	horzcat([])	minus(-)	realpow	tril
atan2	cscd	hsv2rgb	mldivide(\)	realsqrt	triu
atan2d	csch	hypot	mrdivide(/)	rectint	true
atand	ctranspose(')	ifft	mtimes(*)	rem	typecast
atanh	cummax	ifft2	mod	repelem	uint16
besselh	cummin	ifftn	mode	repmat	uint32
besseli	cumprod	imag	NaN	reshape	uint64
besselj	cumsum	Inf	ndims	rgb2hsv	uint8
besselk	deg2rad	int16	ndgrid	rmfield	uminus(-)
bessely	diag	int32	ne(~=)	round	unwrap
beta	diff	int64	nextpow2	sec	uplus(+)
betainc	dot	int8	nnz	secd	vander
betaincinv	double	inv	nonzeros	sech	var
betaln	eig	ipermute	norm	sign	vertcat([;])
bicg	end	isempty	normest	sin	xor
bitand	eps	isequal	not(~)	sind	zeros
bitor	eq(==)	isequaln	nthroot	single	
bitxor	erf	isfinite	num2cell	sinh	
bsxfun	erfc	isfloat	numel	size	
cart2pol	erfcinv	isinf	nzmax	sort	

Processing quite big data

Multiplication of 2 NxN matrices

```
>> C = A * B
```

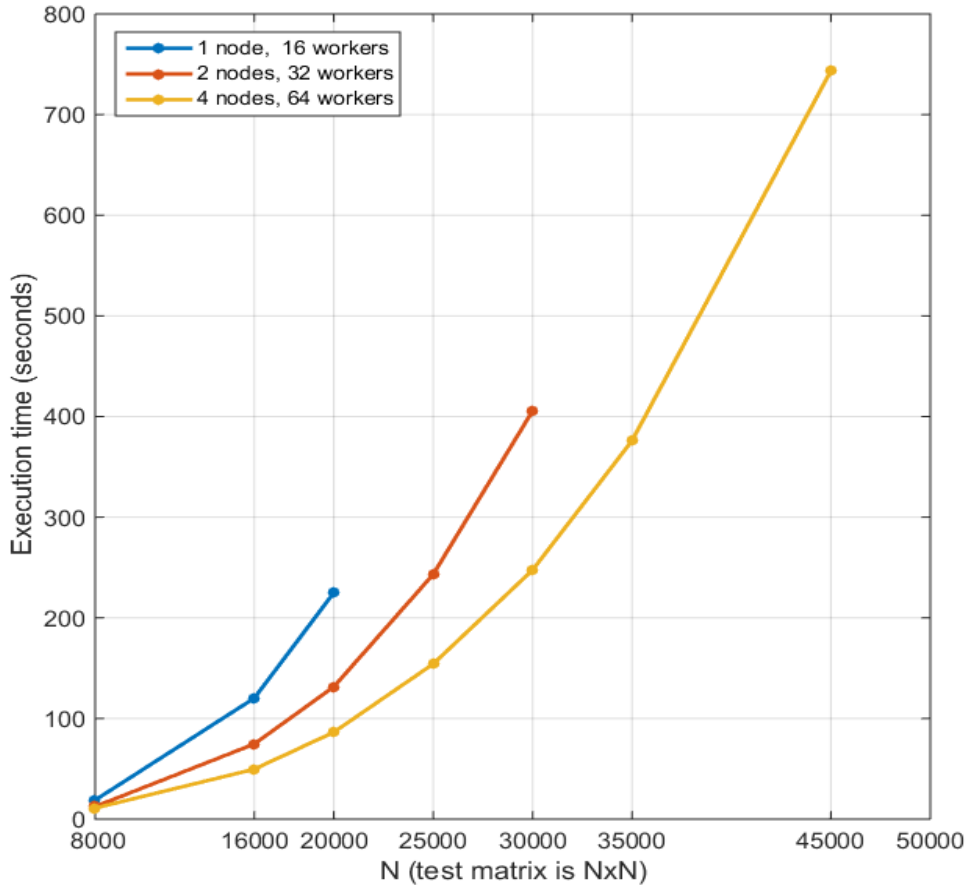

N	Execution time (seconds)		
	1 node, 16 workers	2 nodes, 32 workers	4 nodes, 64 workers
8000	19	13	11
16000	120	75	50
20000	225	132	86
25000	-	243	154
30000	-	406	248
35000	-	-	376
45000	-	-	743
50000	-	-	-

Processor: Intel Xeon E5-class v2
16 cores, 60 GB RAM per compute node, 10 Gb Ethernet

Processing quite big data

Distributed DEMO

Agenda

- How big is big?
- Reading big data
- Processing quite big data
- Processing big data
- Summary

Processing really big data

When you can never see all the data

- Can never have all the data loaded
- Must process small pieces of data independently
- Extract (“map”) some pertinent information from each independent piece
  - Typically summary statistics, example records, etc.
  - No communication between pieces
- Combine (“reduce”) this information to give a final (small) result
  - Intermediate results from each piece must be communicated

Introduction to Map-Reduce

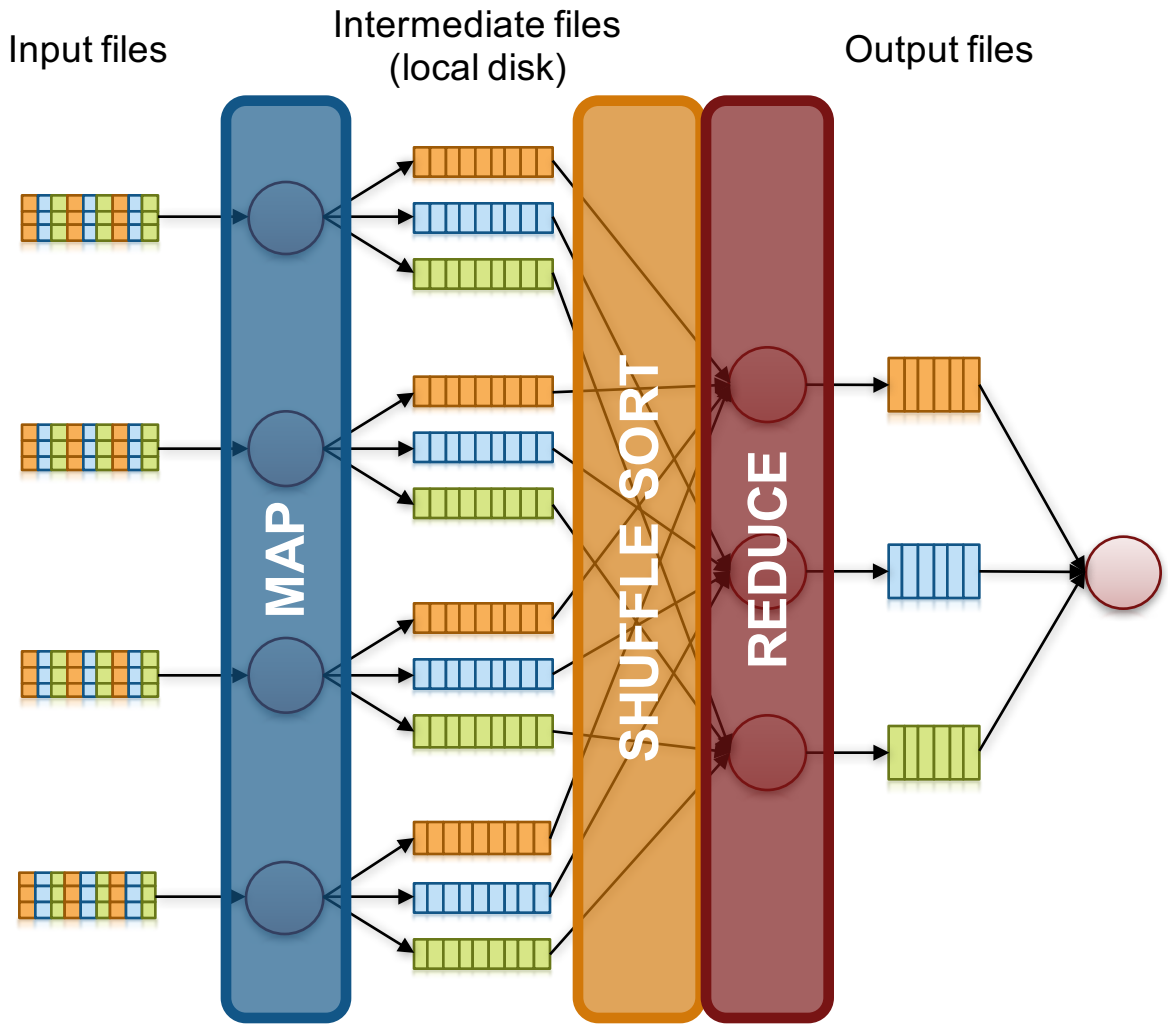

Introduction to Map-Reduce

Example:
National popularity contest

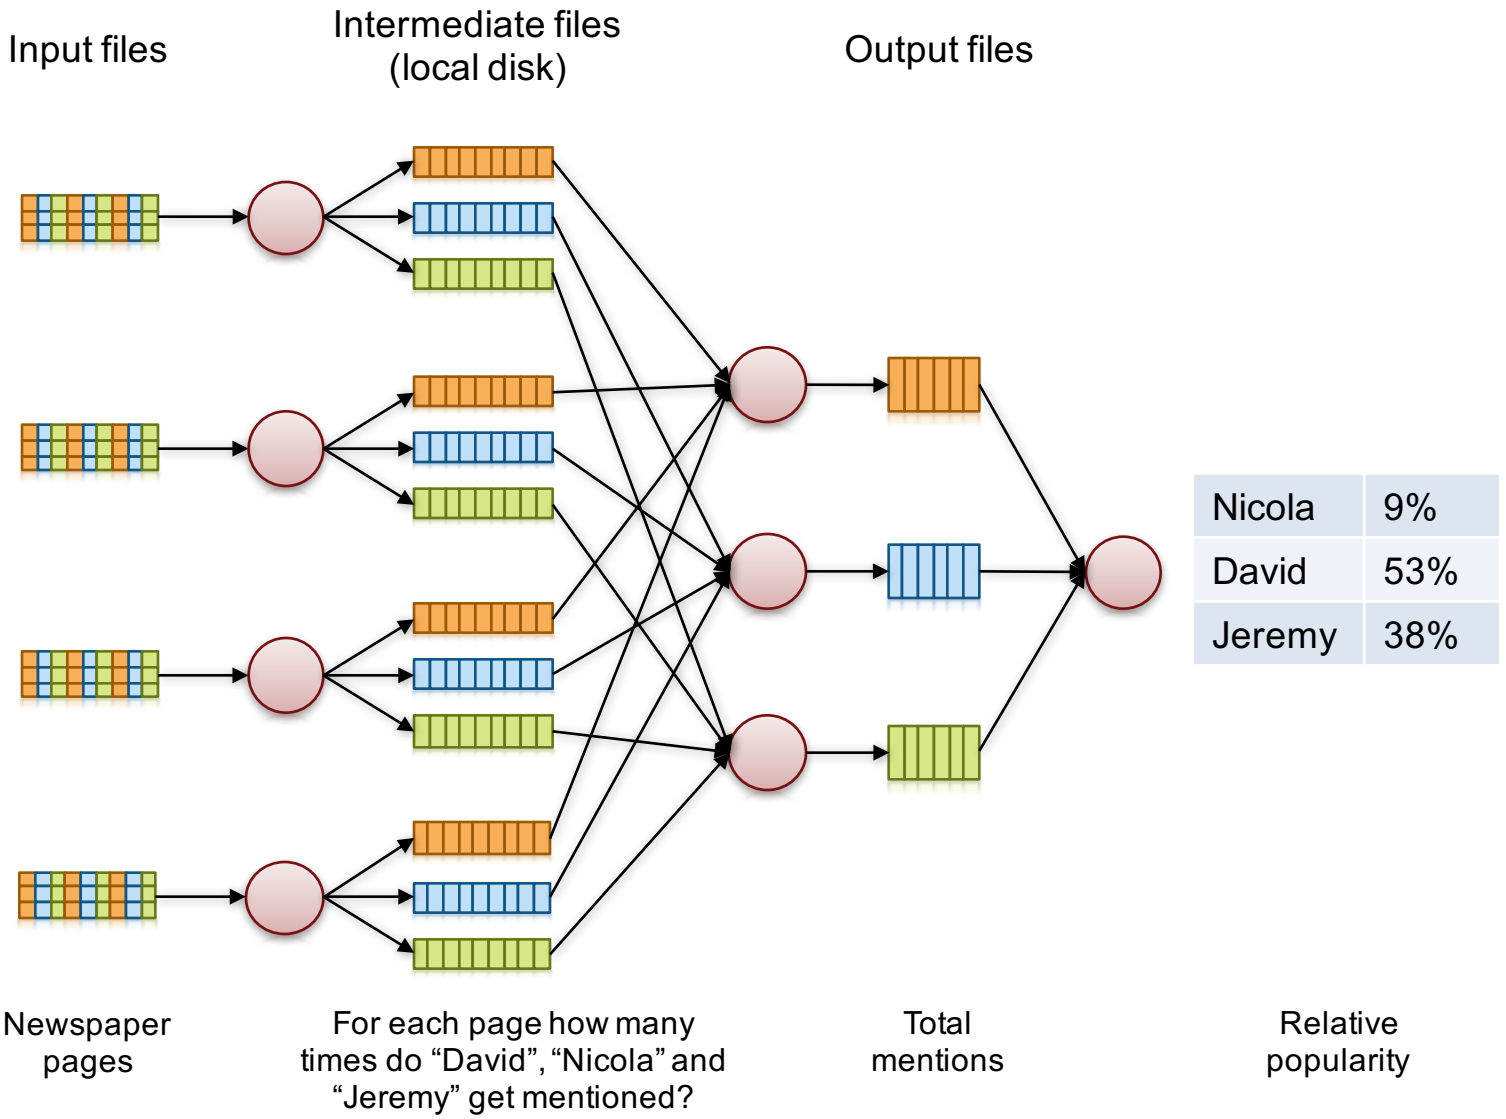

Processing medium data

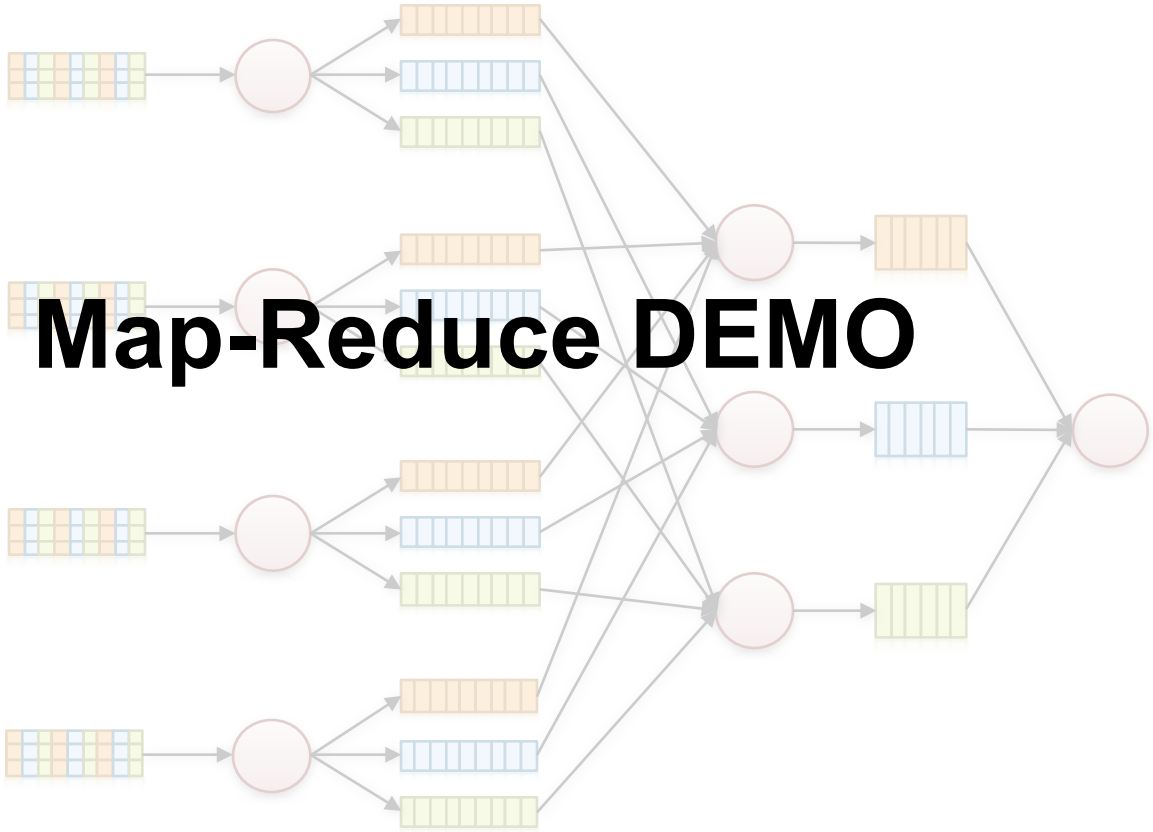

MATLAB

with Hadoop

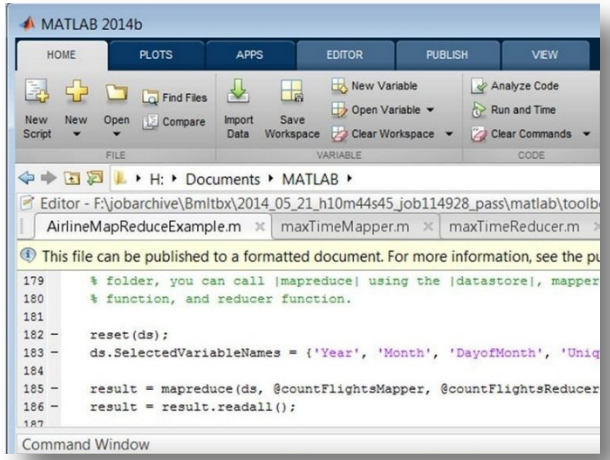

`Datastore` access data stored in HDFS from MATLAB

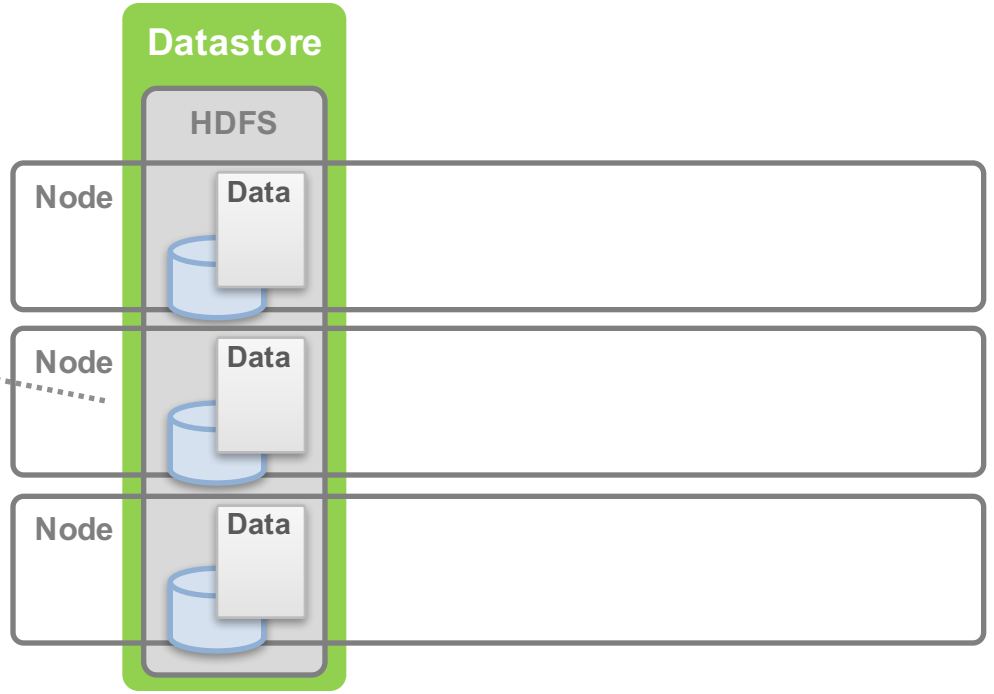

Hadoop

MATLAB Distributed Computing Server

with Hadoop

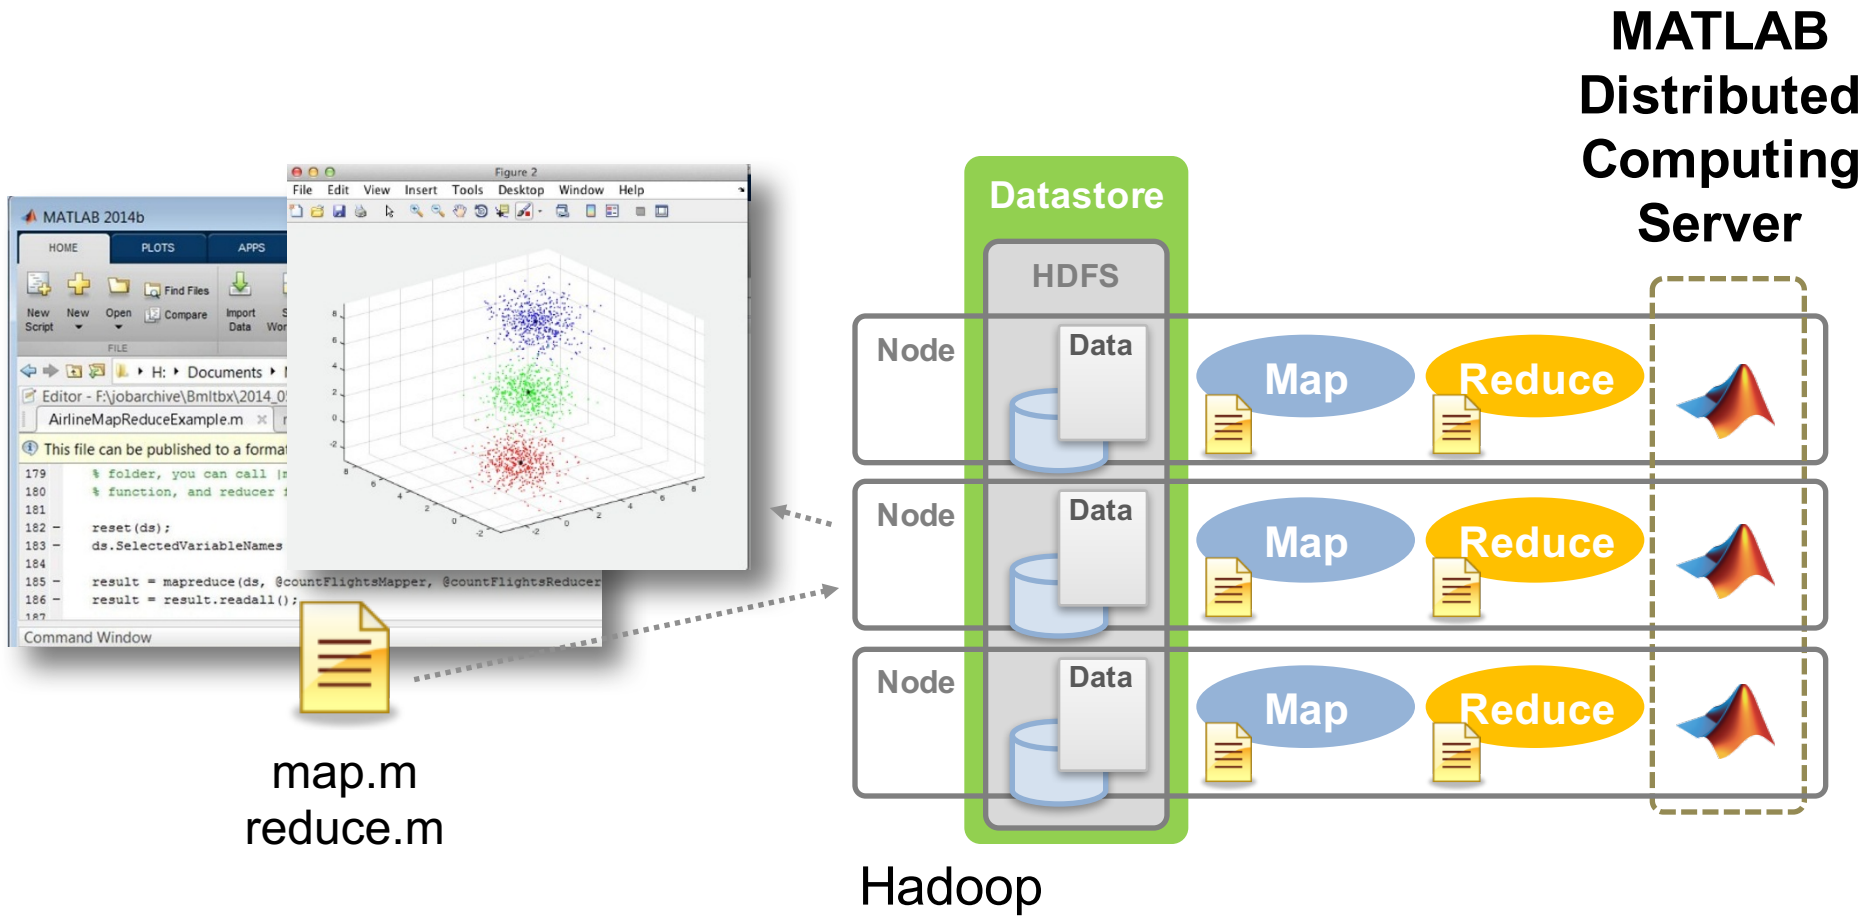

Agenda

- How big is big?
- Reading big data
- Processing quite big data
- Processing big data
- Summary

Techniques for Big Data in MATLAB

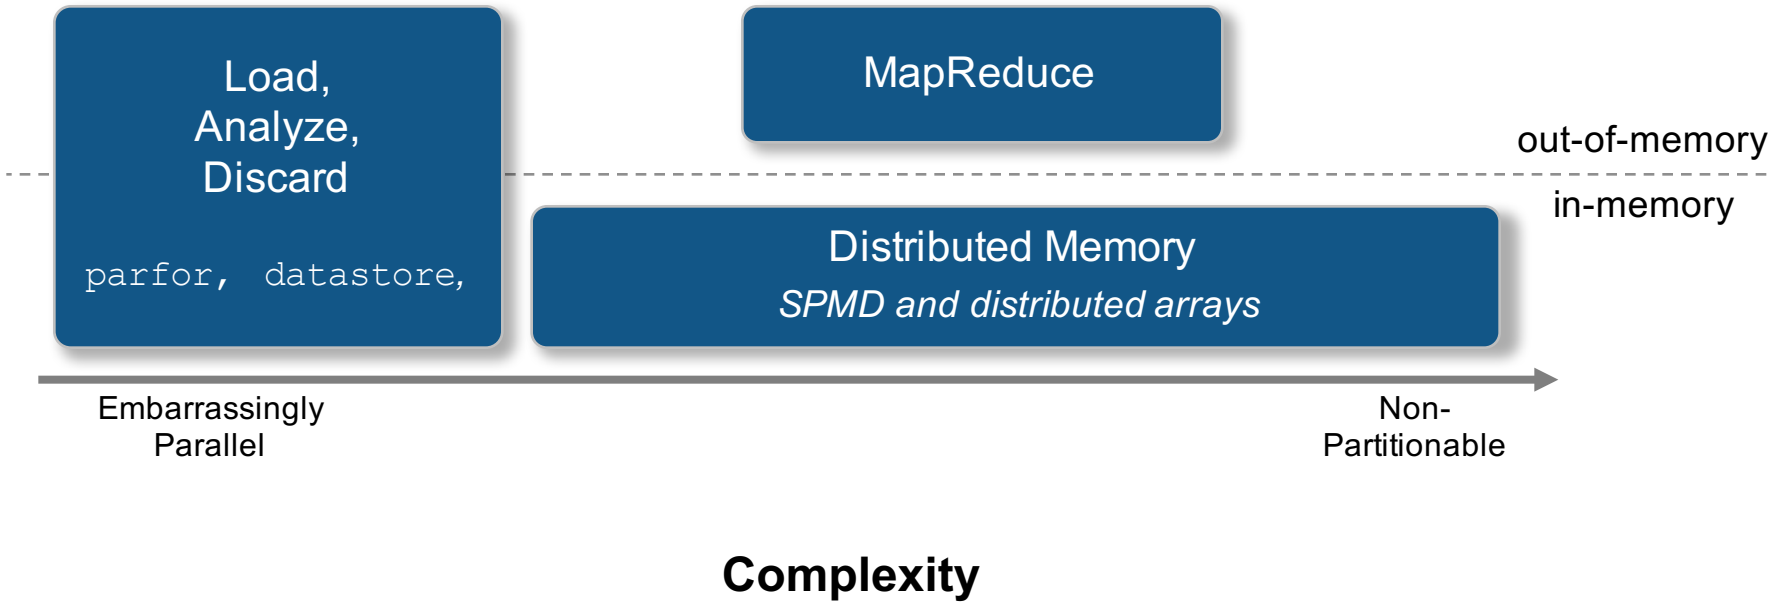

Summary

Reading data:

1. When data gets big you need to work on pieces
2. Use datastore to read pieces from a large data-set

Processing data:

1. If it fits in memory, use MATLAB
2. If it fits in cluster memory, use distributed arrays
3. If you need to scale beyond cluster memory, use map-reduce