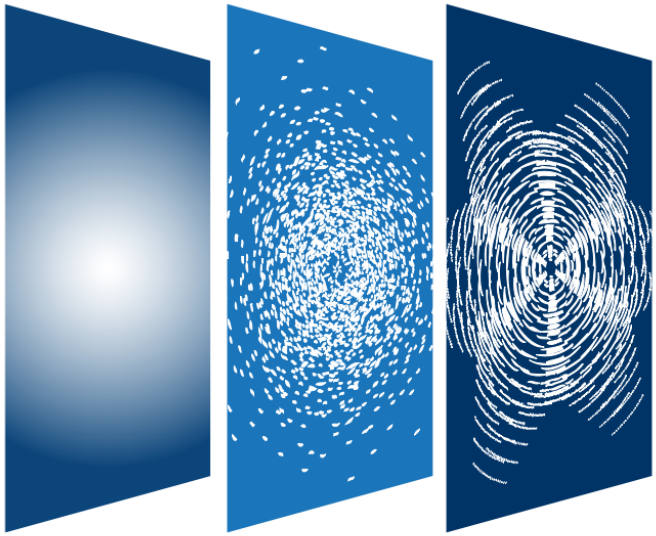

CASA

Common Astronomy Software Applications

CASA Performance and Status

Jeff Kern

CASA Group Lead

James Robnett

Science Computing Infrastructure Lead

Atacama Large Millimeter/submillimeter Array

Expanded Very Large Array

Robert C. Byrd Green Bank Telescope

Very Long Baseline Array

CASA Performance

- Realized CASA Performance is a function of several factors:
  - Usage Patterns
  - Data Access Rate (Disk I/O)
  - Memory Utilization
  - Processor Speed

- CASA team is addressing bottlenecks:
  - Plotting (particularly CalTables)

We have optimized CASA for modern data from VLA and ALMA

e.g. We assume the visibility data does not fit in system memory

Memory

- CASA assumes that the visibility data is too large to hold in memory.
  - True for modern ALMA and VLA data sets
  - Implies that imaging is the dominant stage in terms of memory usage

- CASA attempts to hold only the necessary information consistent with efficient operation in memory.
  - For basic imaging (no A- or W-terms, single scale, single frequency term) the peak memory usage is approximately 24 bytes / pixel
    - 1 Double Precision Complex Grid
    - 1 Single Precision Complex Grid
  - For spectral line this is per output channel

Image Size	Memory Required
3k x 3k	206 MB
6k x 6k	864 MB
10k x 10k	2.4 GB

Memory

- MS-MFS can be the limiting step in memory usage:
  - Convolution Functions (in Images): $\frac{1}{2}(N_S^2 N_T^2)(20 S_{Max}/ N_{pix})^2$
  - Residual Images: $N_S N_T$
  - Model Images: N_T
  - Cache Functions: Approximately 30%

N_S	N_T	S_{Max}/ N_{pix}	Images	3k x 3k	10k x 10k
5	2	1.5 %	45	1.6 GB	18 GB
5	2	2.5%	52	1.8 GB	20.8 GB
5	2	5%	89	3.2 GB	35.6 GB

Data Access (Disk I/O)

- Although some stages of processing access only a fraction of the data, many operations require accessing the full visibility data
  - For many operations the I/O access either in terms of absolute data rate or number of transactions is the limiting factor in CASA performance.
- CASA is working on improvements which decrease the data access requirements, and ensuring that reads are as efficient as possible.

	Peak Data Rate	OS Realized Data Rate	Time for 500 GB file
SATA Disk	115 MB/s	60 MB/s	2h 22 m
Raid	200-500 MB/s	375 MB/s	22 m
Lustre (10 GB)	1.2 GB/s	900 MB/s	9.5 m

Data Access (Disk I/O)

- Unlike many other aspects of the processing equation, time alone will not solve the data access problem.

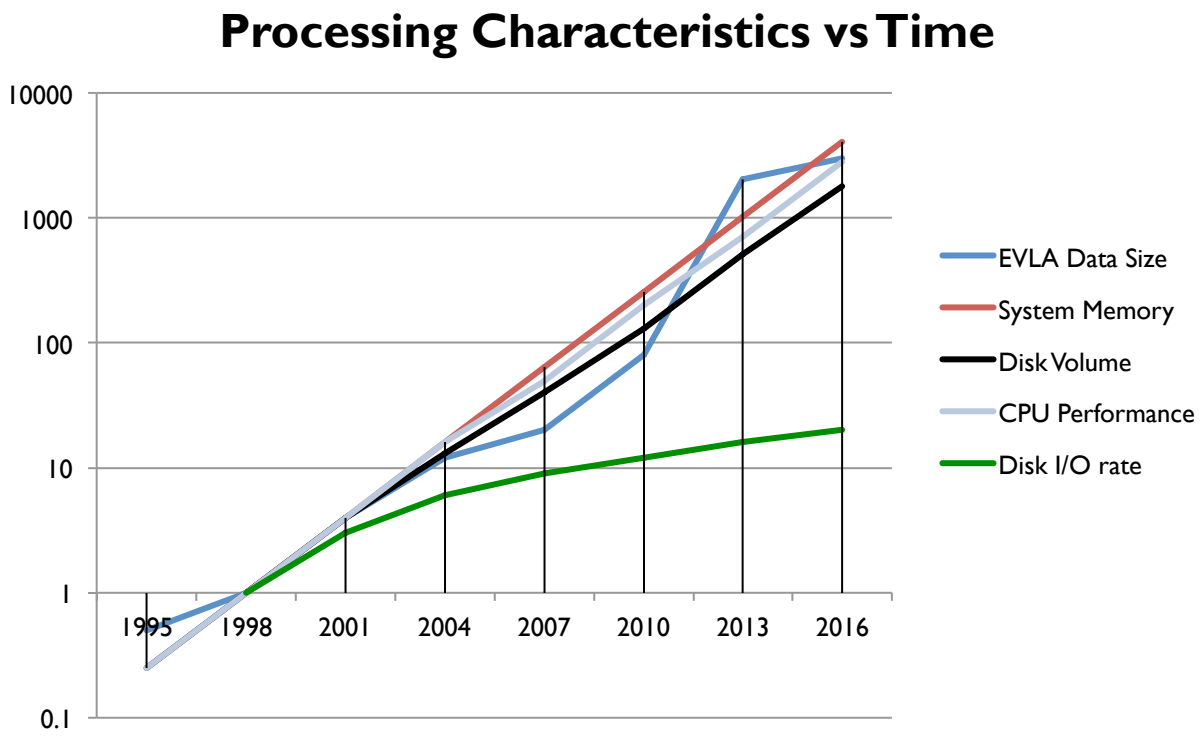

Processor Speed

- Processor clock speeds are no longer dramatically increasing
  - Industry has gone to more cores, rather than faster clocks
- Successive generations of processors will show modest increases in performance due to improved instruction pipelines etc.
  - Only weakly affects observed performance of serial CASA

Currently CASA is more limited on data access and memory footprint than on processor speed. We are currently trading increased computation load for decreased memory and data access.

Processor Speed

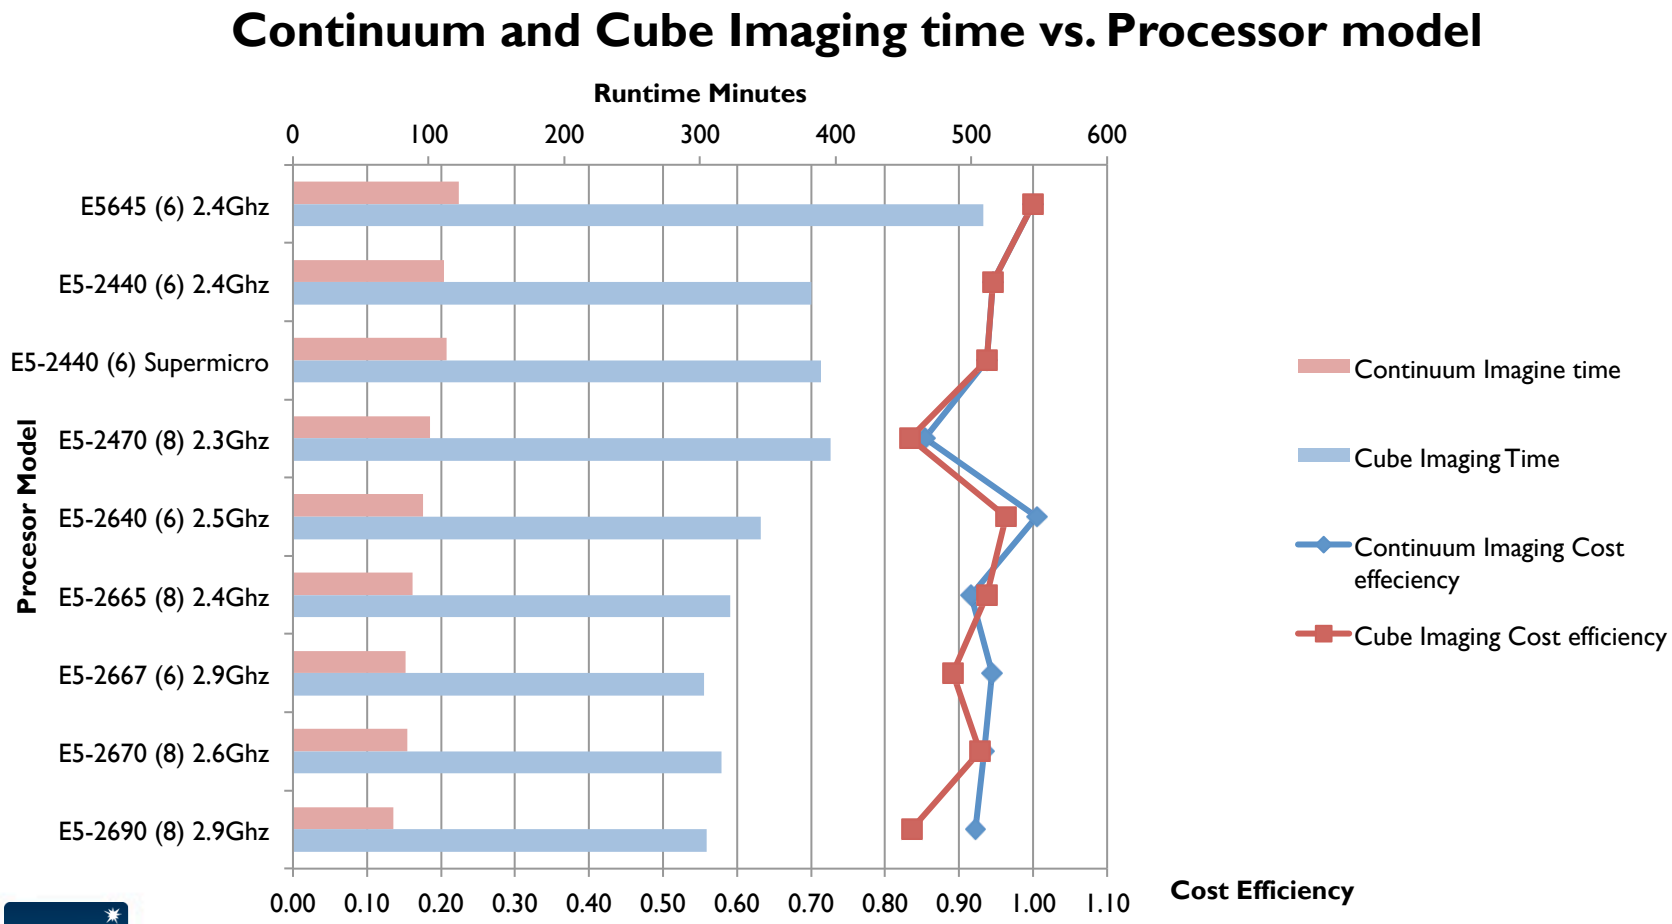

So What Should I Buy?

- Well… it depends:
  - What type of data do you expect?
  - How will you process it?
  - How long are you willing to wait for processes to complete?
  - How much do you want to spend?
  - How much IT support do you have?

- CASA will continue to evolve, we are continually working to bring algorithms to the community which optimize our hardware utilization.
  - This means the correct answer will change over time as we respond to changes in algorithms and hardware ecosystem.

That doesn't help!

https://science.nrao.edu/facilities/vla/data-processing/hardware-recommendations

General Guidelines

- Do not neglect I/O
  - Multi Disk RAID is probably the best option for small groups.
  - Large Groups w/ IT support might consider High Performance File Systems
- Consider your memory requirements.
  - 16 cores doesn't help if you can only run 1 engine because of memory.

Update pending (Coming in early May)

Where we put our Money

- Dual Intel E5-2670 (2.6GHz) 8 core processors
- 64GB of memory
- 40Gbit/s Infiniband NIC
- 875 TB Lustre System
  - 20 GB/s
  - $250K
- Total Cost per Node: $6K
  - Not Including Storage

Labor Not Included!

Parallel CASA

- Many stages of process are embarrassingly parallel.
  - CASA has the tools to support parallel processing
  - Not yet a completely turn-key solution
- Some members of this workshop have been using parallel portions of the system

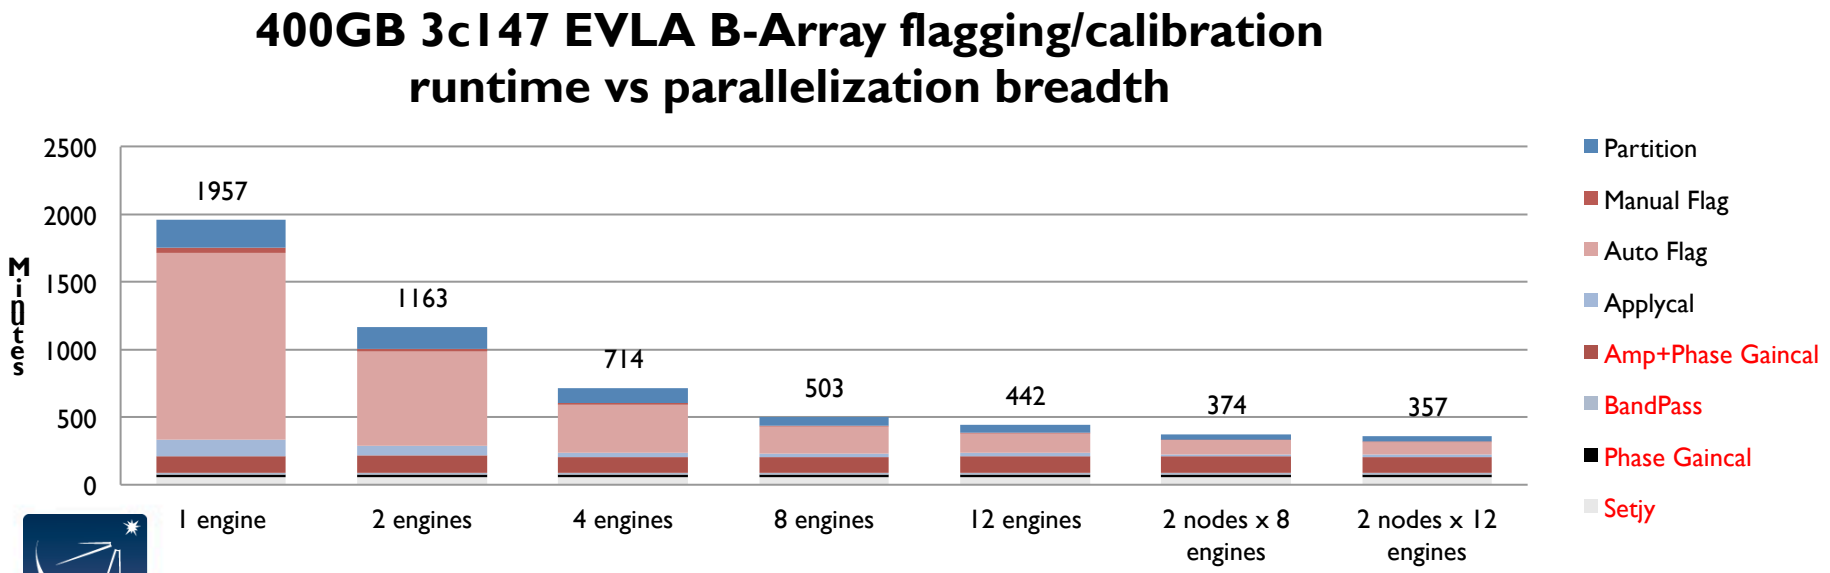

Parallel CASA

- The technique is to create a Multi-MS (MMS):
  - Eventually this will be completely transparent to Users

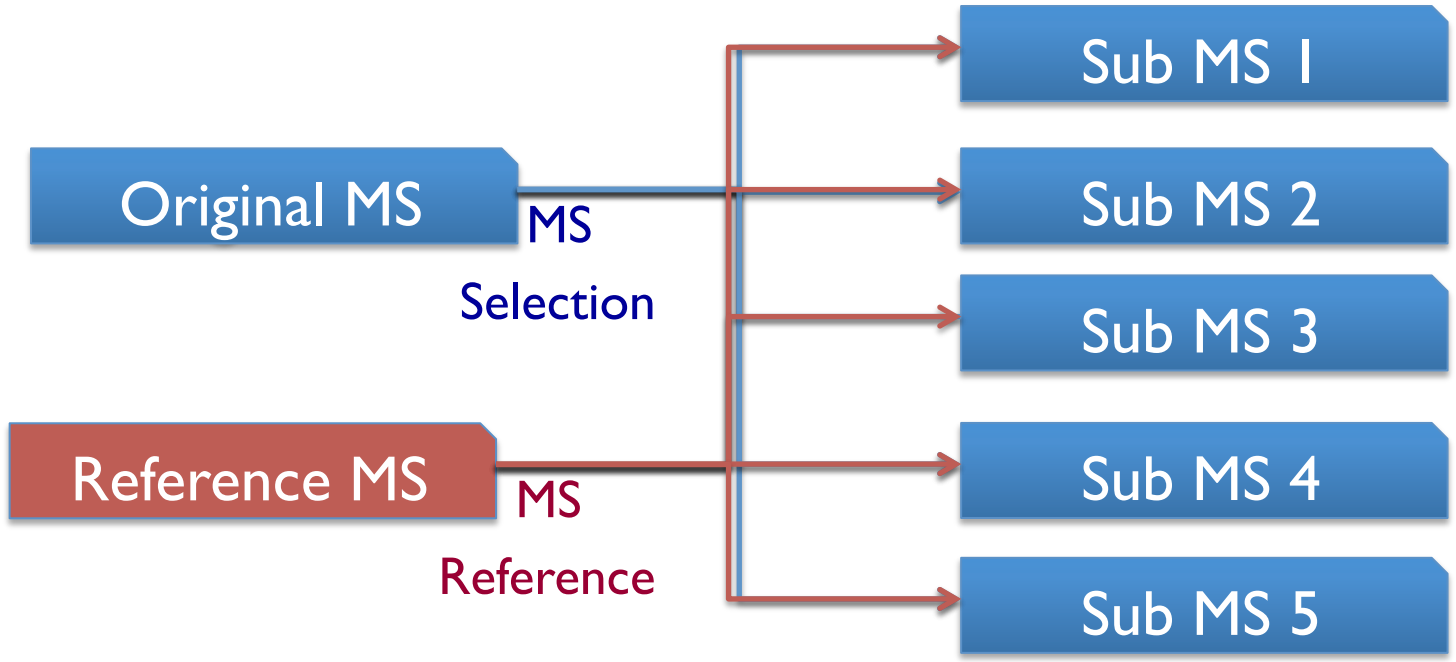

- Supported by partition (4.0) and mstransform (4.1) tasks

Parallel CASA

- Most time consuming tasks which do not produce a new MS are now MMS aware and will automatically parallelize where possible.
  - By default we use 80% of the available cores on the current host
  - Usage can be configured using cluster spec file.

- Currently tasks that write a new MS usually write a non MMS out
  - With CASA 4.1 the *mstransform* task makes most of these tasks parallel capable

- Beginning to parallelize the pipelines
  - Ensure "standard path" works

Parallel CASA: Imaging

- Parallel Imaging is implemented in the task *pclean*
- Continuum
  - Technique is to segment the data, gridding each portion separately and then combine for the minor cycle.
    - Recall that this is the memory intensive portion
  - If you don't have enough memory, you can thread the gridding up to 4 ways (1,2, or 4 threads are supported)
    - Not quite as efficient, but less memory intensive
  - MS-MFS and Mosaic are not currently supported
- Spectral Line
  - Segment data on output channel, fully parallel imaging
  - Stitch cube back together at the end

Currently the parallel and serial versions of imaging use different code. This can lead to slightly different behaviors / results.

Parallel *CASA: Continuum Imaging*

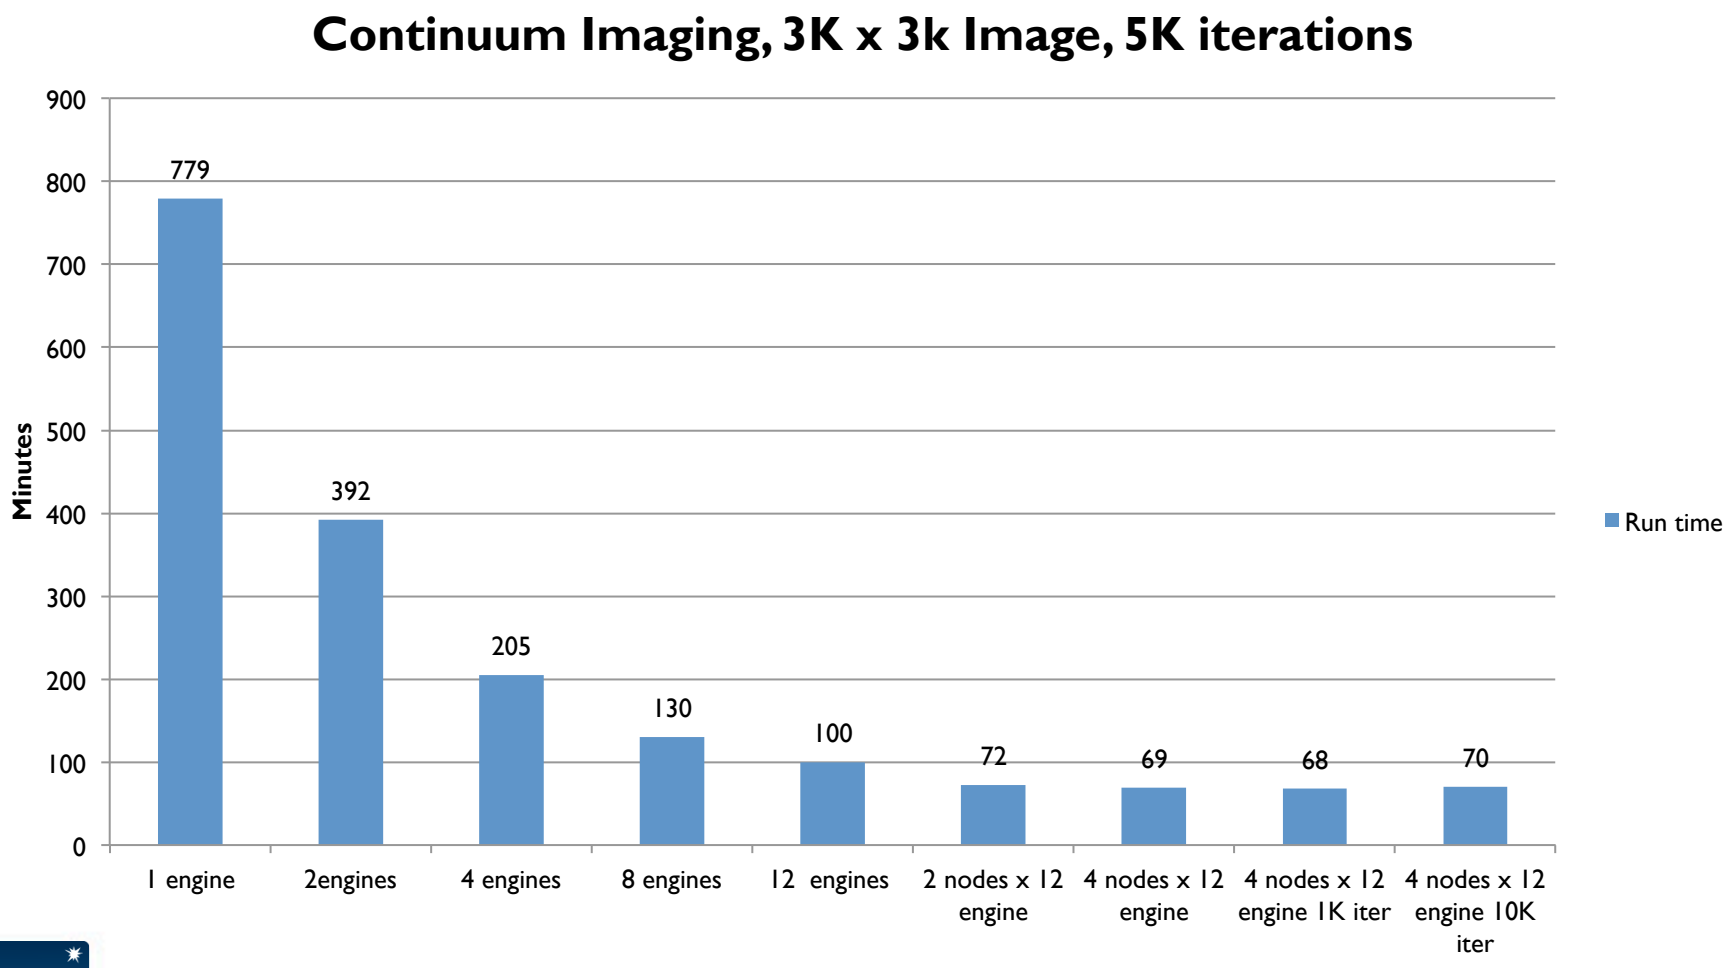

Parallel CASA: Spectral Imaging

Spectral Line Imaging, 3K x 3K image x 1K channels, 5000 iterations

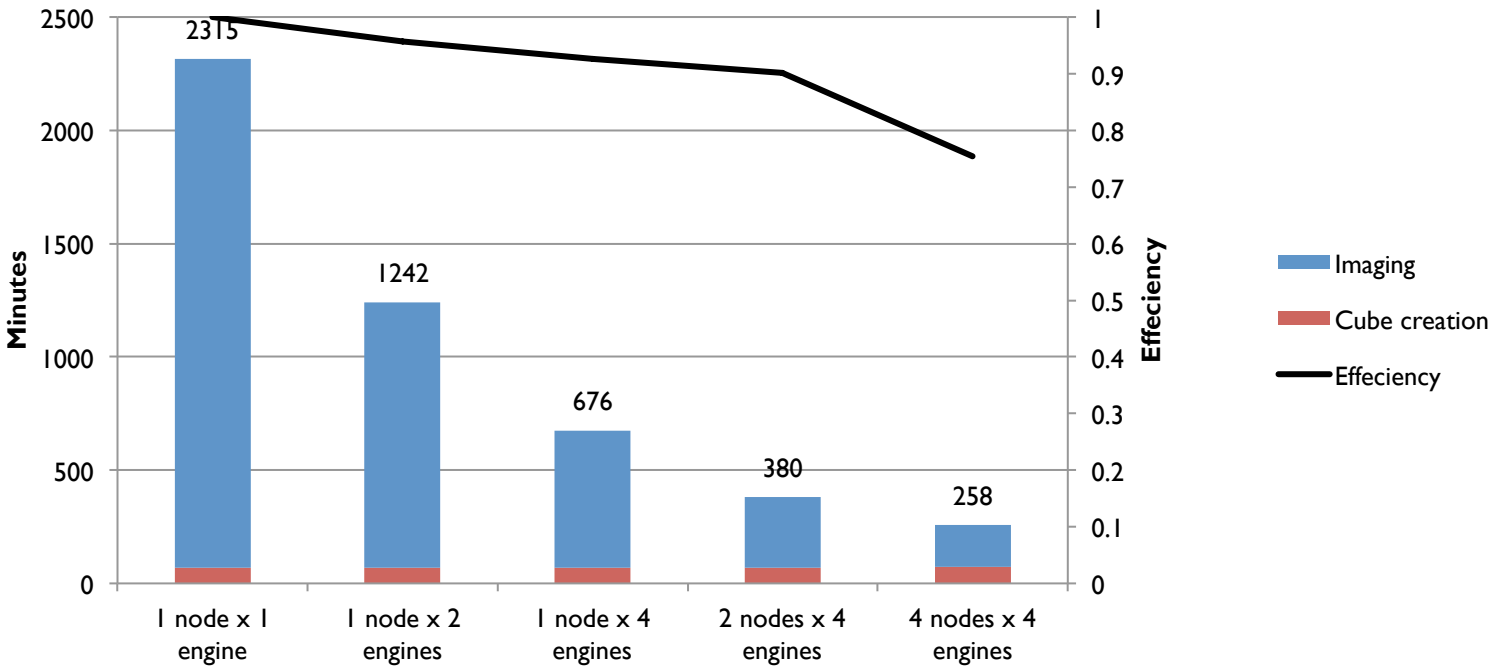

New Features in 4.1

- CASA 4.1 is expected out end of this month
- Viewer:
  - Histogram Support
  - Spectral Line Fitting
  - PV Diagrams
- Flagging
  - Flagdata on Cal Tables
  - Multiple input lists
- Listobs Improvement
  - msmd tool
- Calibration
  - Improved support for ObsId
  - *Imaging Sensitivity Calculator*
  - PlotMS
    - Cal Table Support
    - *Multi-Panel Display*
  - mstransform task
    - *Still in early testing*
  - Single Dish
    - Modified for Cal Tables
    - Improve BL subtraction
  - Imaging
    - *pclean support for msmfs*

Longer Term

- CASA is continuing to evolve:
  - *plotms* rework: batch mode, wider application
  - Cal Library
    - Removes the *applycal* step, decreases disk use, less IO
  - *mstransform* task
    - Decreases passes through data
  - Data weights
    - Correct handling of spectral weights
  - Imager Re-Architecting
    - Same code for parallel and non-parallel clean
    - Increased flexibility in options
  - VO access in the viewer

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

www.nrao.edu • science.nrao.edu