Development and status of early pipelines for MWA and PAPER

Daniel C. Jacobs
Arizona State University

MWA
MURCHISON
WIDEFIELD
ARRAY

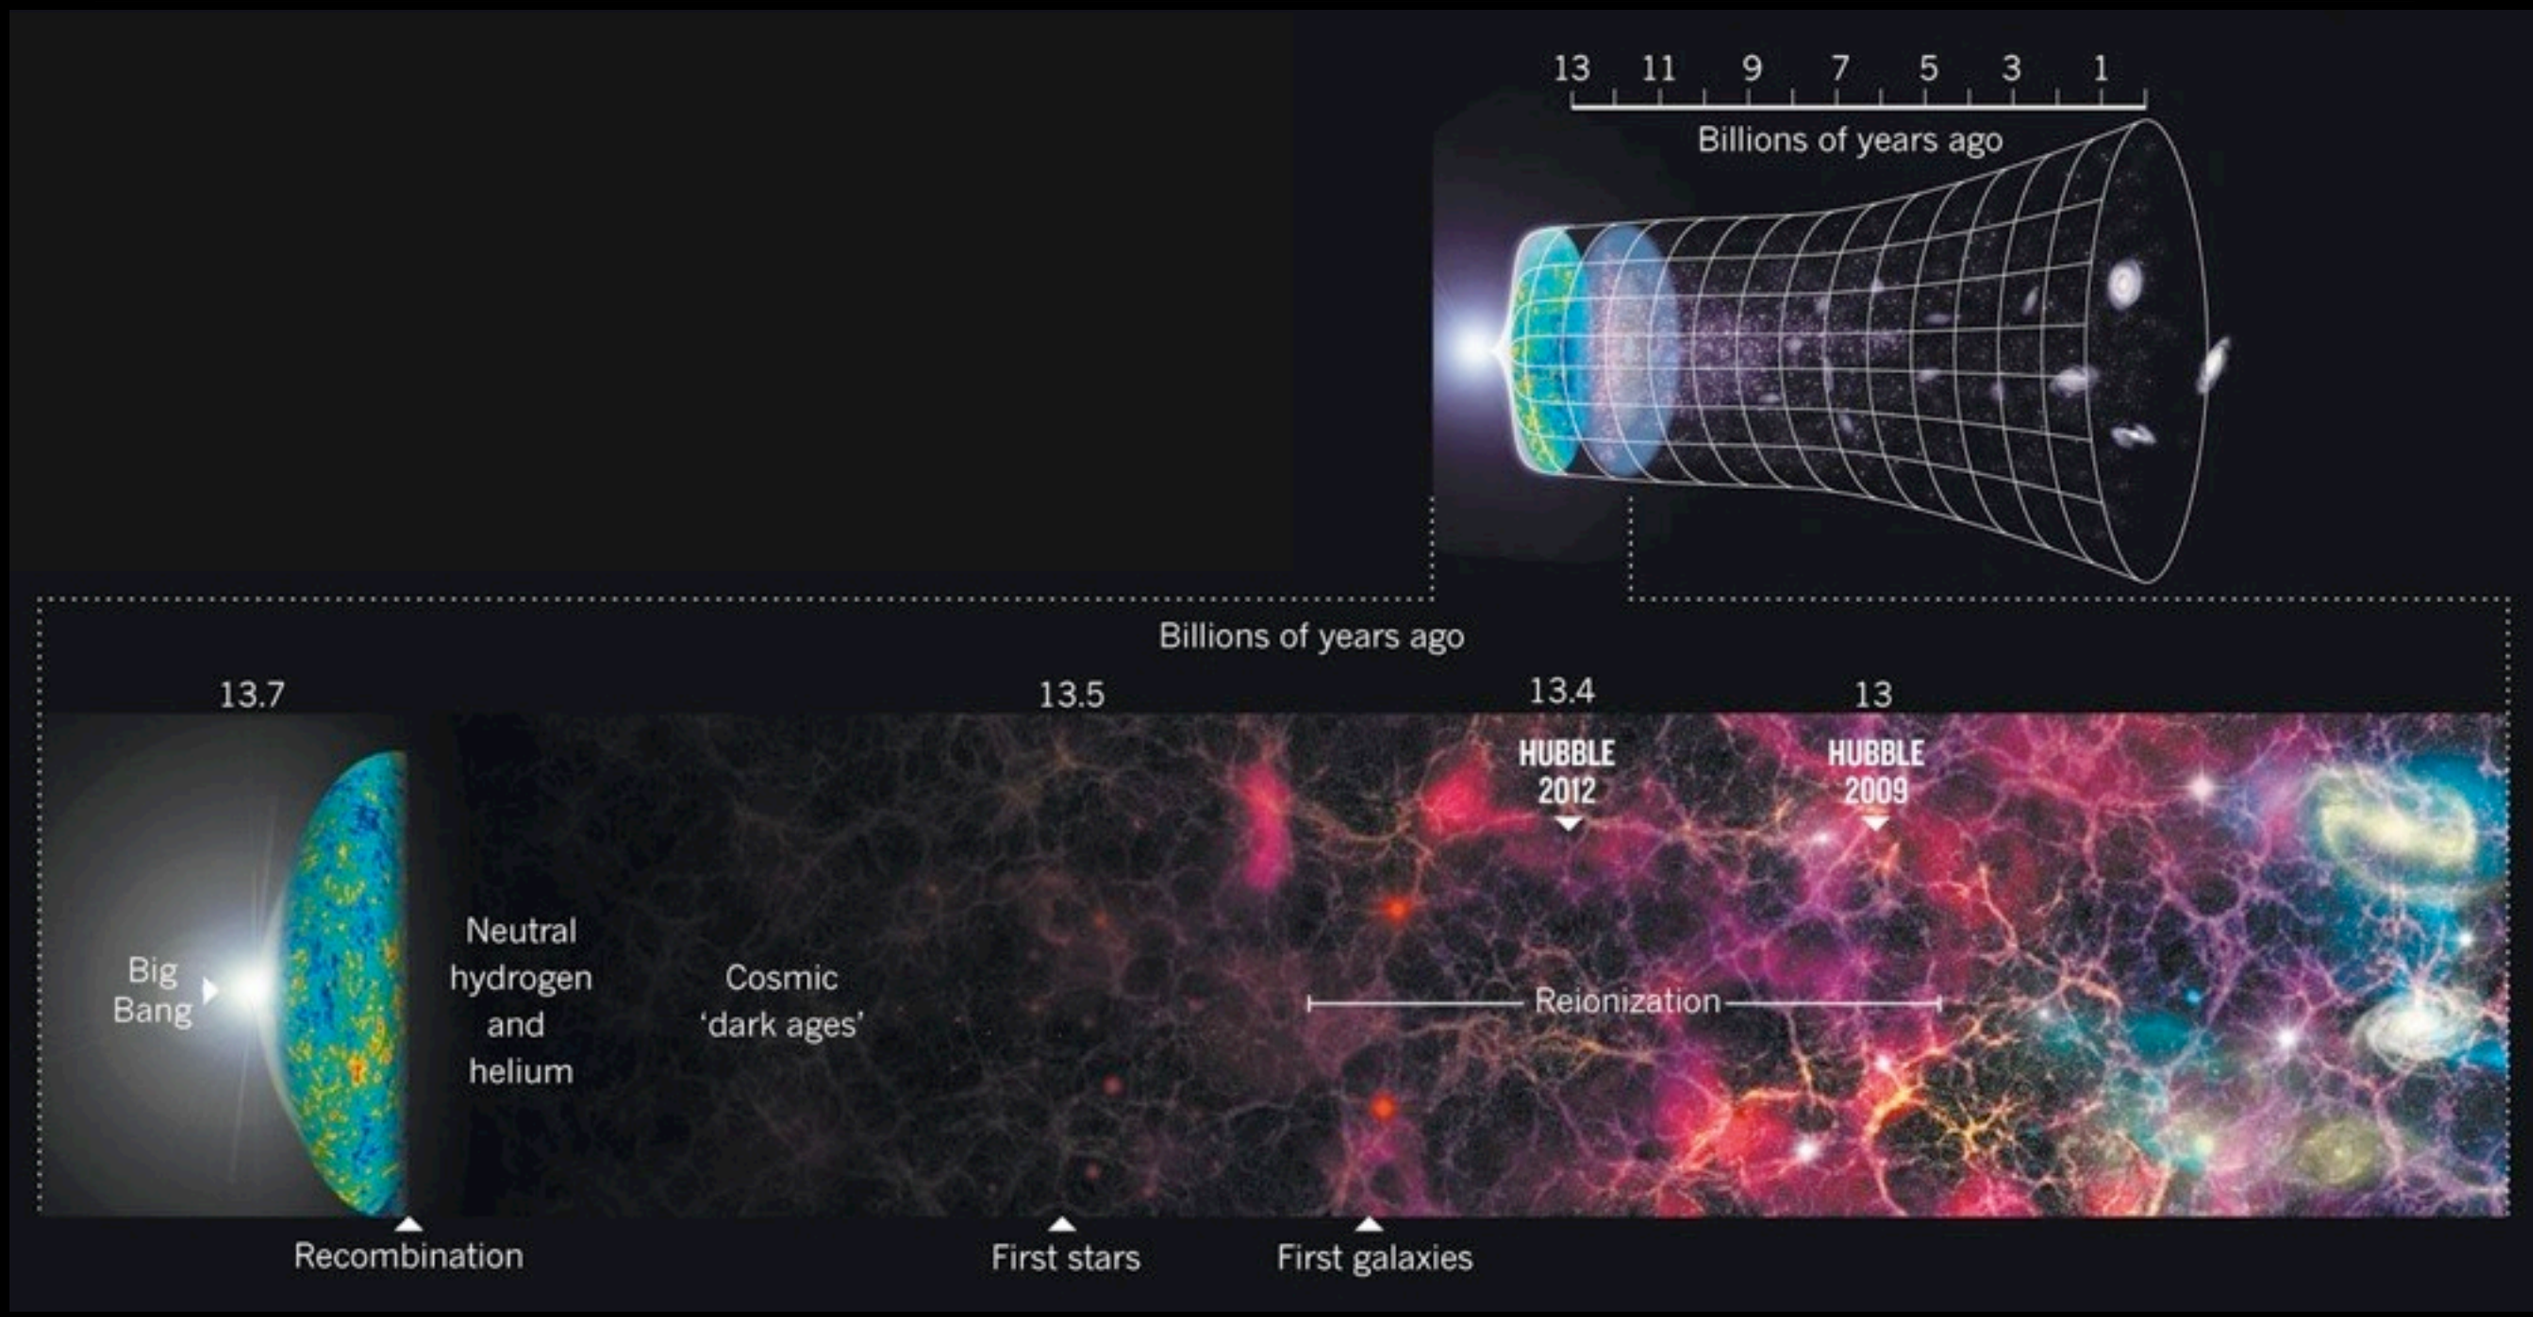
13
11
9
7
5
3
1
Billions of years ago
Billions of years ago
13.7
13.5
13.4
13
HUBBLE 2012
HUBBLE 2009
Big Bang
Neutral hydrogen and helium
Cosmic 'dark ages'
Reionization
Recombination
First stars
First galaxies

Cowen 2013

Foregrounds

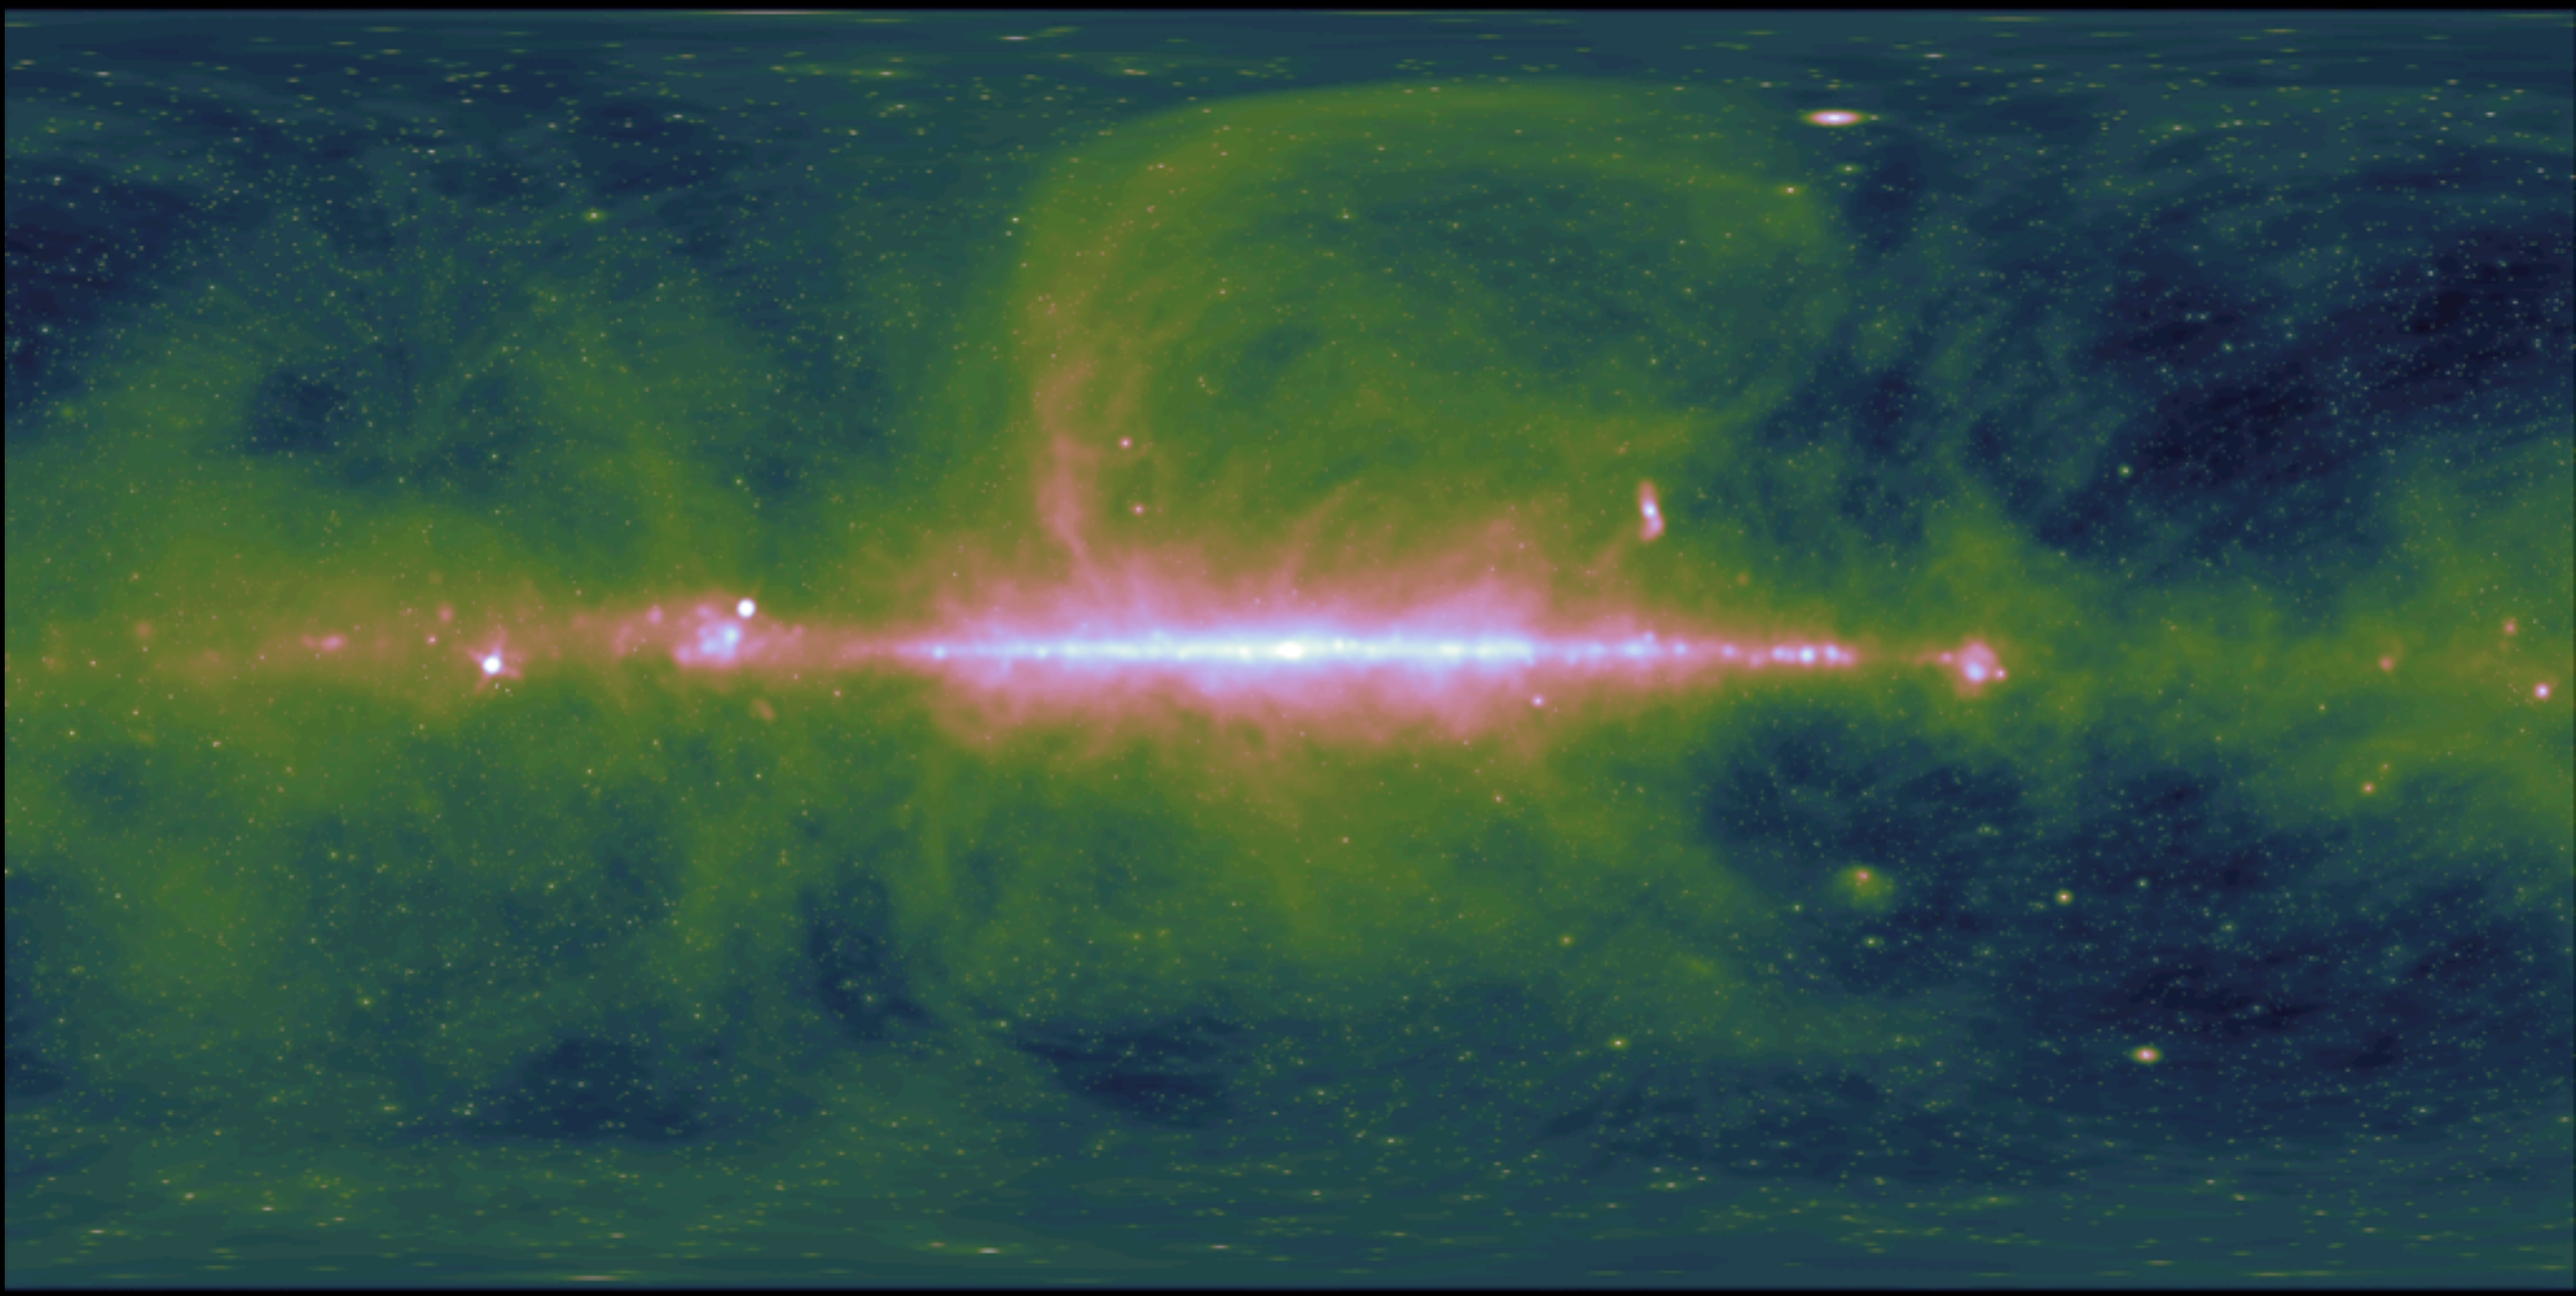

http://loco.lab.asu.edu/danny_jacobs/
google: danny radio sky poster

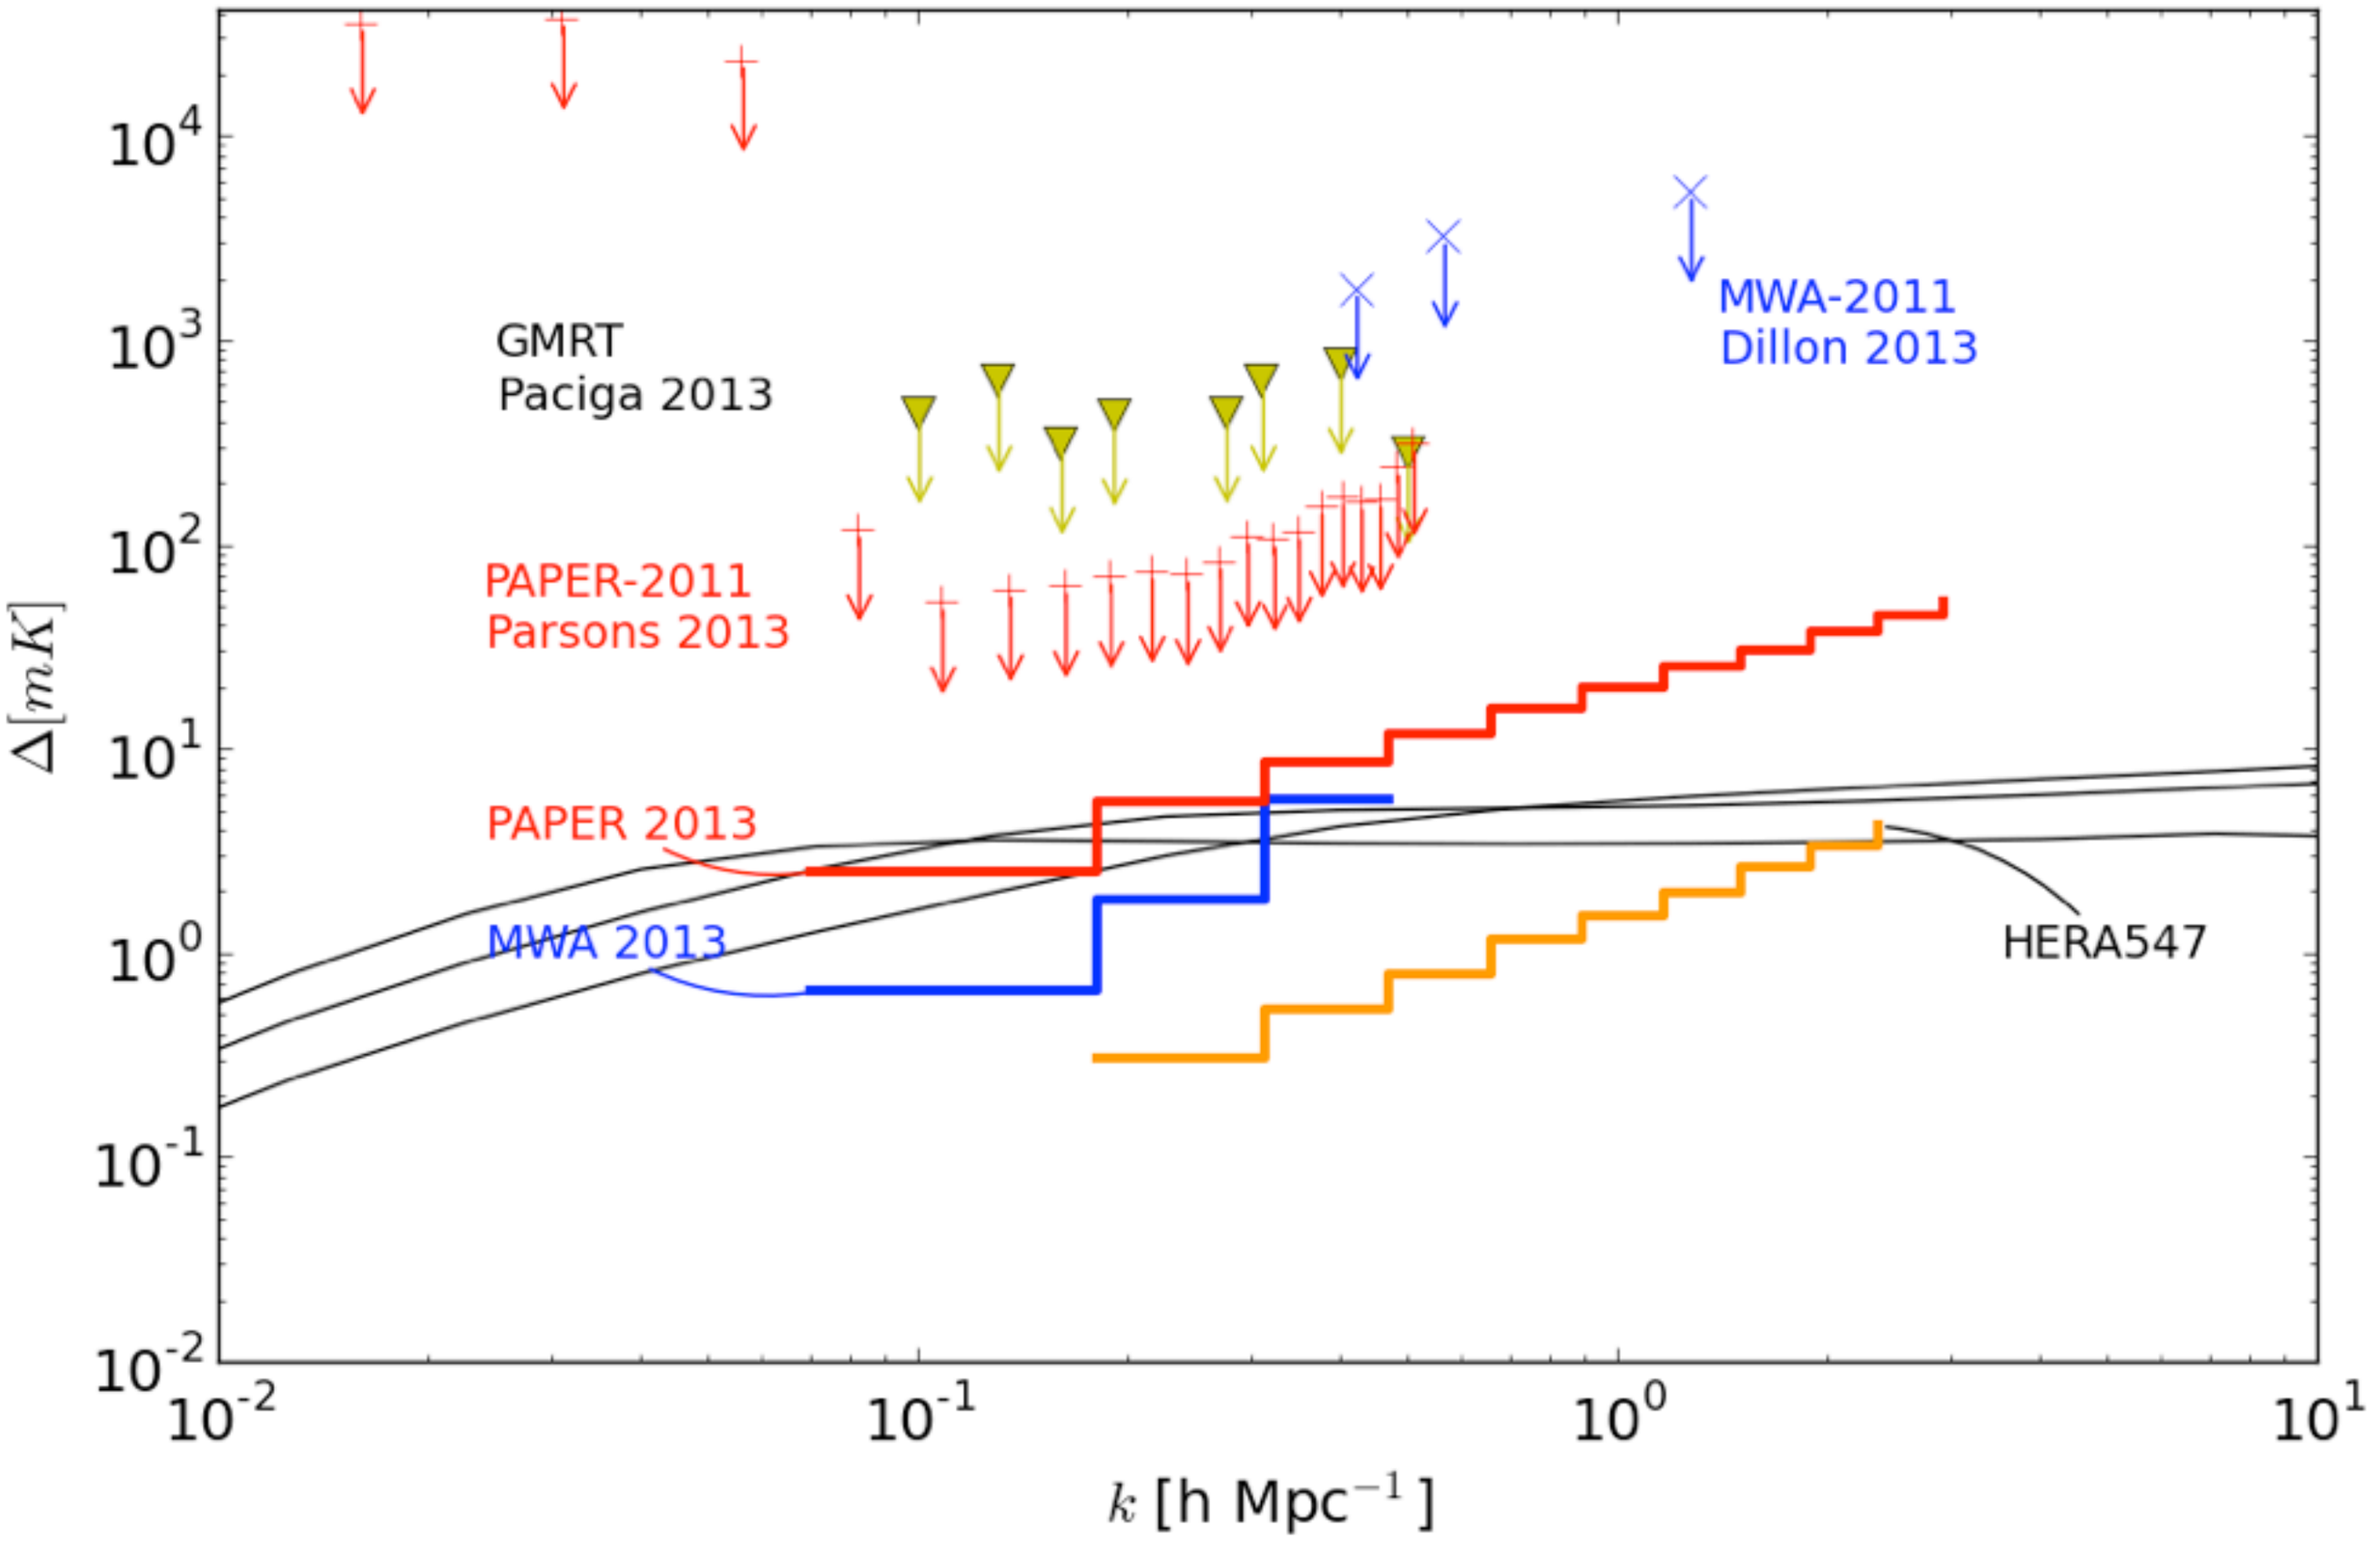
GMRT
Paciga 2013
PAPER-2011
Parsons 2013
MWA-2011
Dillon 2013
PAPER 2013
MWA 2013
HERA547
$\Delta[mK]$
k [h Mpc^{-1}]
10^{4}
10^{3}
10^{2}
10^{1}
10^{0}
10^{-1}
10^{-2}
10^{-2}
10^{-1}
10^{0}
10^{1}

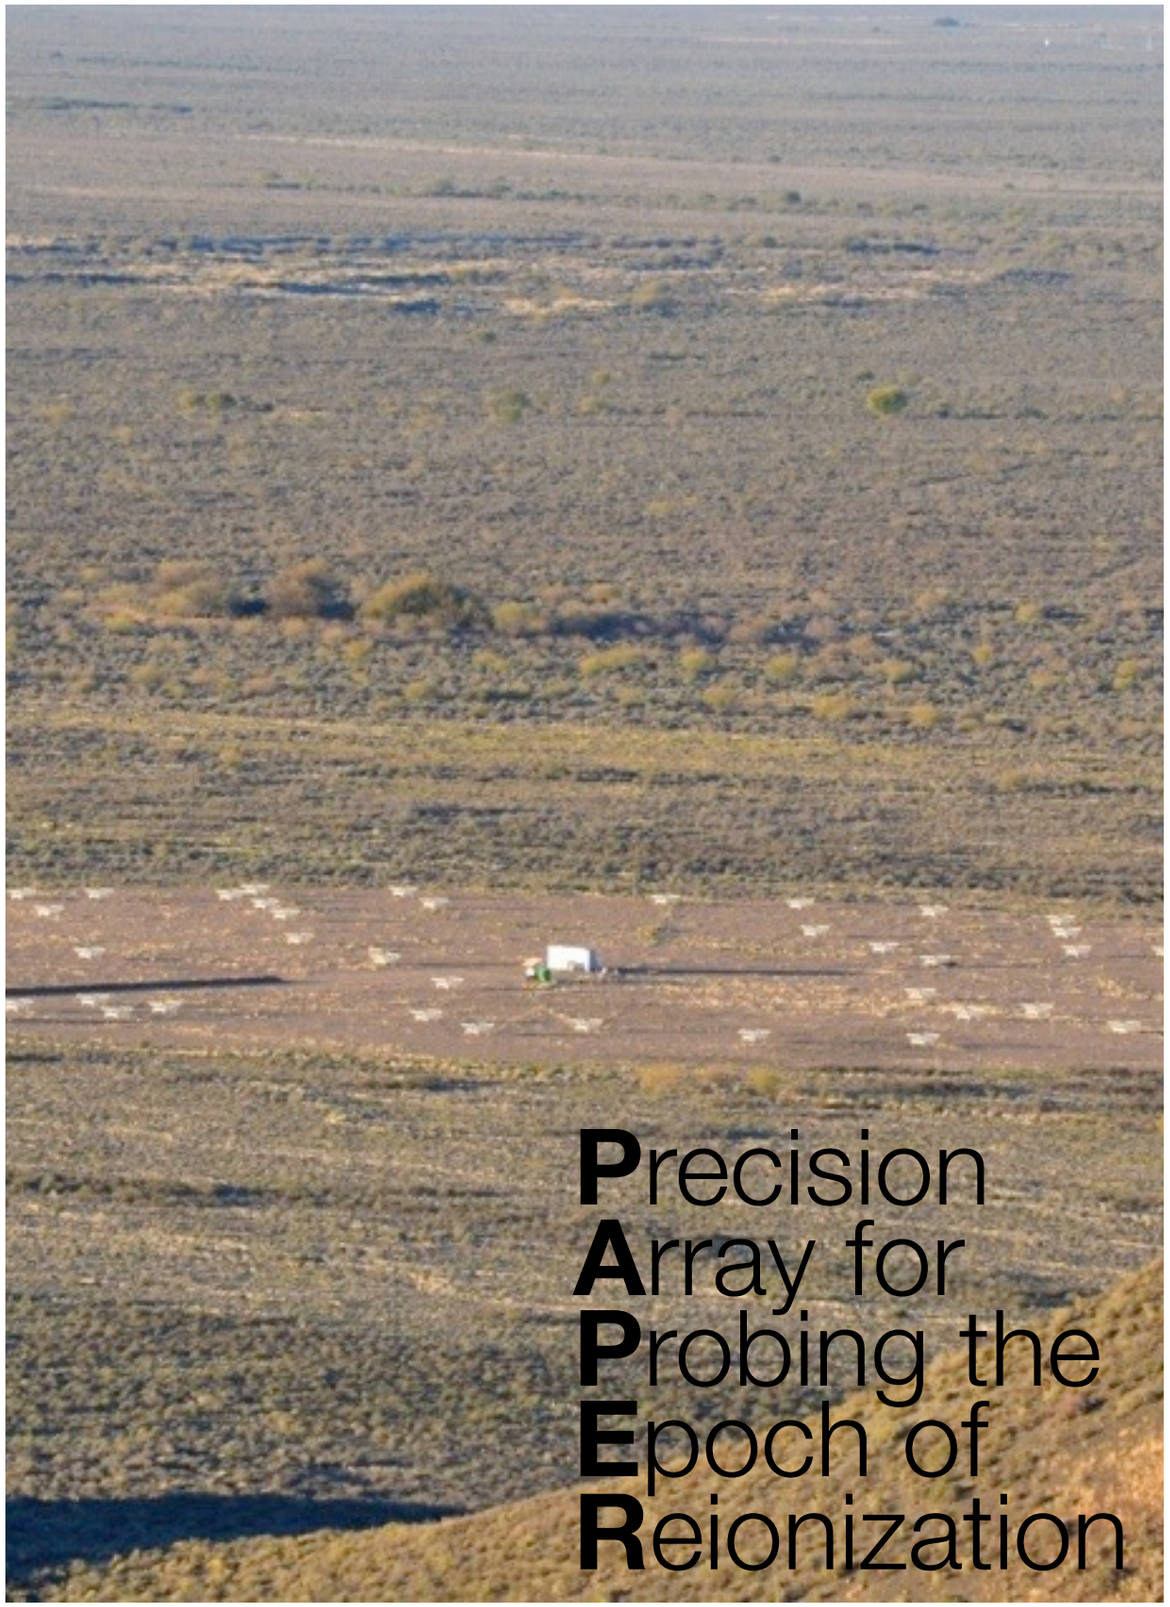

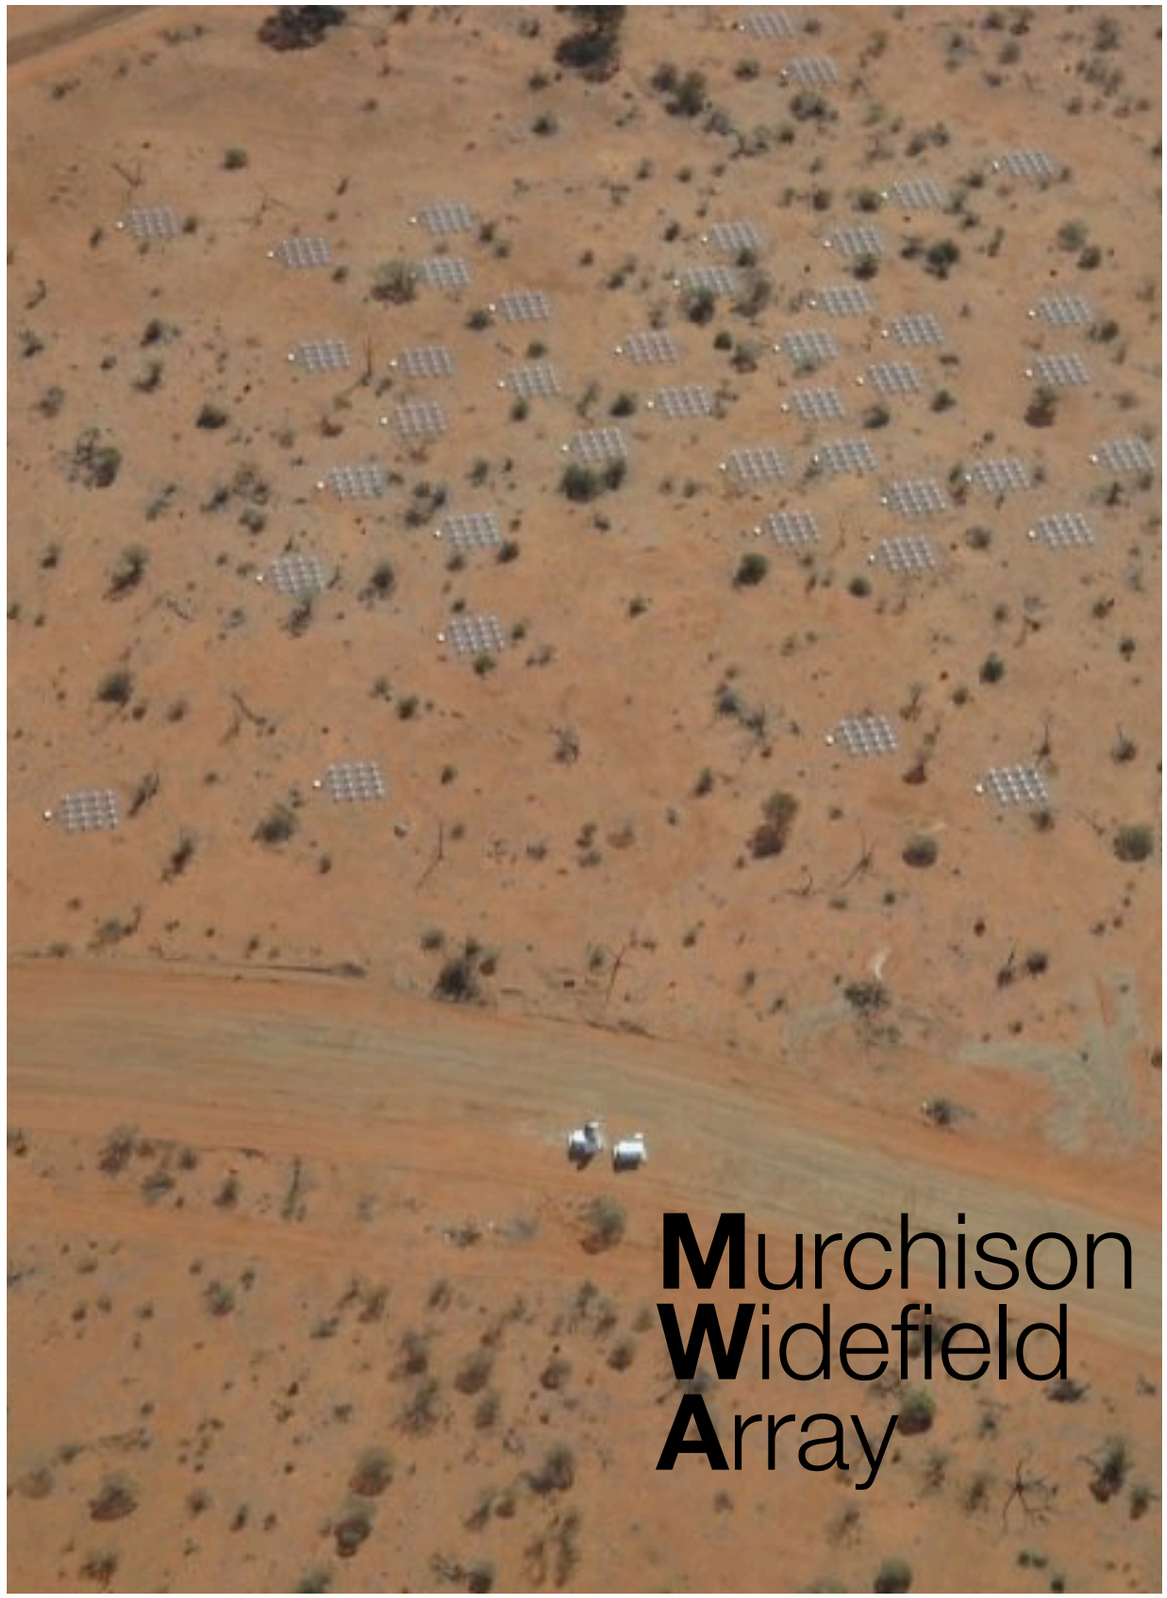

MWA
MURCHISON
WIDEFIELD
ARRAY

Radio Frequency Interference

PAPER Test Bed
Green Bank, WV

PAPER
Karoo South Africa

MWA
Murchison Shire, WA

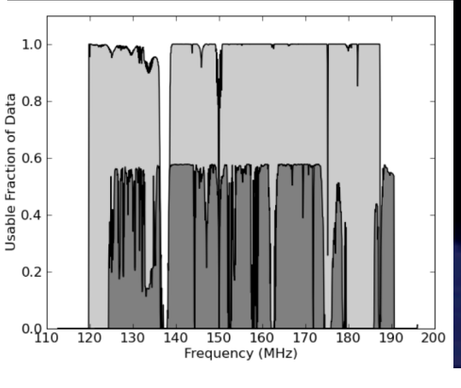

Human Interference

PAPER Test Bed
Green Bank, WV

PAPER
Karoo South Africa

MWA
Murchison Shire, WA

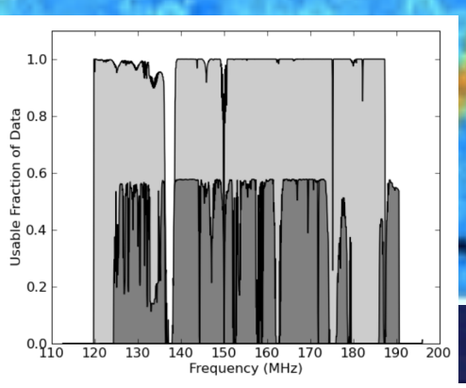

MWA
MURCHISON
WIDEFIELD
ARRAY

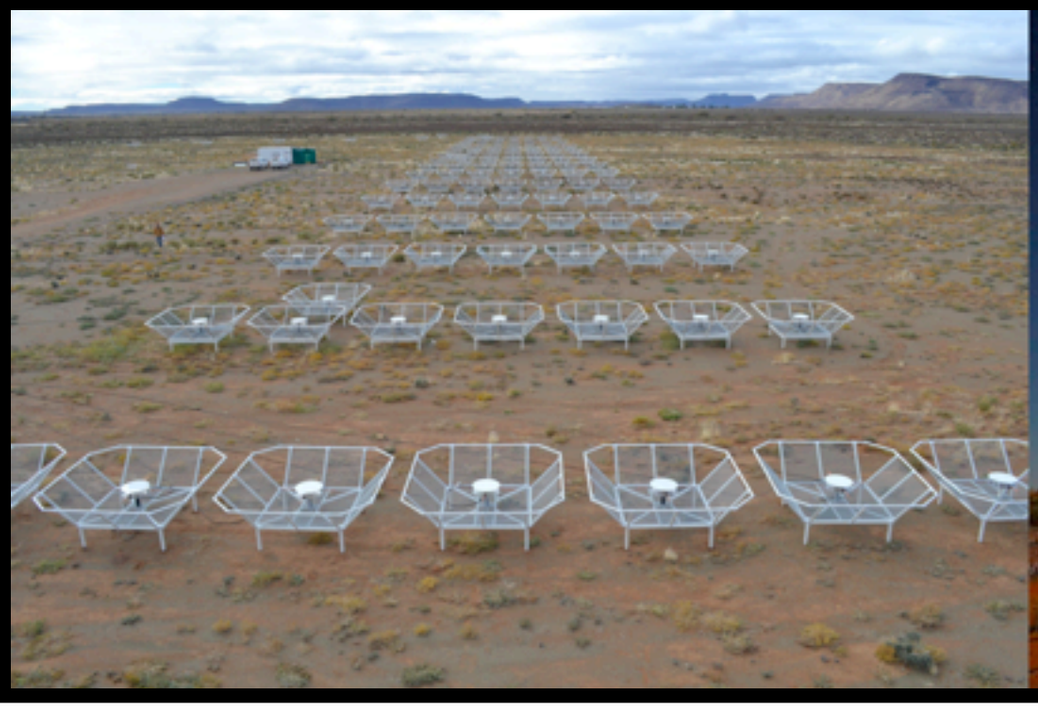

	PAPER	MWA
Antenna	dipole	phased dipole tiles
antenna positions	grid	radial
spectrum	100-200MHz	80-300MHz
location	Karoo Desert (SKA-South Africa)	Western Australia (SKA-Australia)
field of view	60 degrees	30 degrees
Strength	systematic rejection	imaging capability
Weakness	limited sensitivity	uneven spectral coverage

MWA and PAPER

MWA
MURCHISON
WIDEFIELD
ARRAY

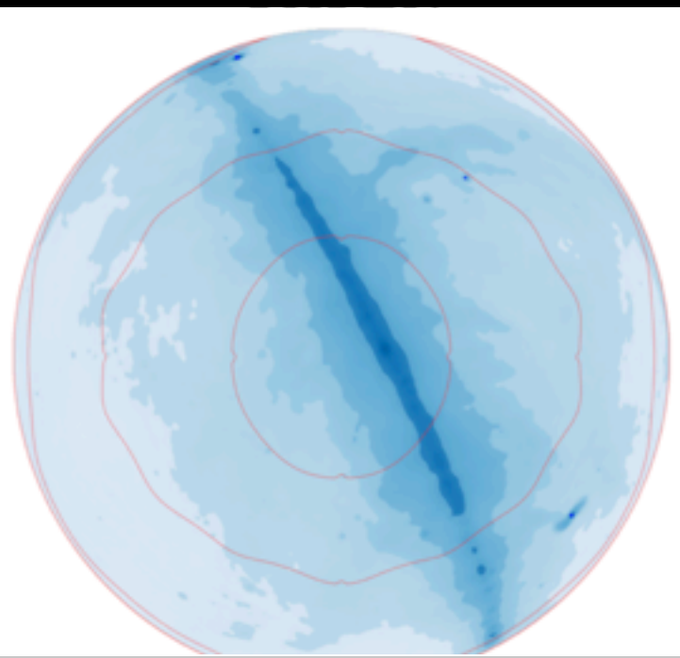

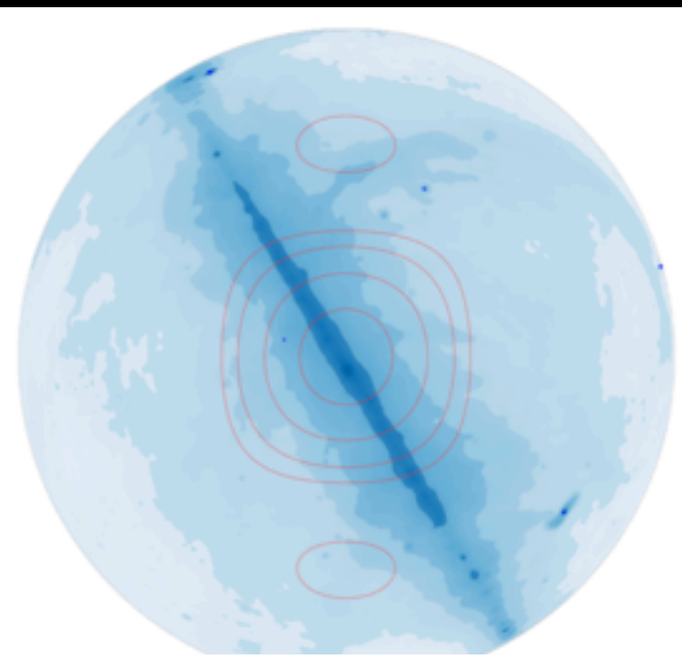

	PAPER	MWA
Antenna	dipole	phased dipole tiles
antenna positions	grid	radial
spectrum	100-200MHz	80-300MHz
location	Karoo Desert (SKA-South Africa)	Western Australia (SKA-Australia)
field of view	60 degrees	30 degrees
Strength	systematic rejection	imaging capability
Weakness	limited sensitivity	uneven spectral coverage

MWA and PAPER

Epoch of Reionization Requirements

100MHz	Foreground removal must have same noise level as reionization signal simultaneous with ~10MHz cosmological bandwidth
100kHz	oversample reionization signal by x10
200m	Maximum useful baseline length sample the brightest modes
10s	preserve fringes (to ≈ 10°) on longest baseline

As built numbers

As built numbers

	PAPER	MWA	HERA

As built numbers

	PAPER	MWA	HERA
Data Rate	25MBps	225MBps	166MBps

As built numbers

	PAPER	MWA	HERA
Data Rate	25MBps	225MBps	166MBps
Bandwidth	10s, 100MHz, 100kHz	2s 30Mhz 20kHz	10s, 100MHz 100kHz

As built numbers

	PAPER	MWA	HERA
Data Rate	25MBps	225MBps	166MBps
Bandwidth	10s, 100MHz, 100kHz	2s 30Mhz 20kHz	10s, 100MHz 100kHz
Total Lifetime Volume	170TB	1.5PB	1.11PB

Some observations about Reionization Data

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

  - the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
    Examples: of things EoR is **not** like:

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)
  - simulations

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)
  - simulations
  - pattern search (spam, intelligence, SETI)

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)
  - simulations
  - pattern search (spam, intelligence, SETI)

- Examples of things EoR IS like:

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)
  - simulations
  - pattern search (spam, intelligence, SETI)

- Examples of things EoR IS like:
  - Deep space communications (super extreme compression)

Some observations about Reionization Data

- on early telescopes it is *information poor*, low SNR, or very high entropy
  - this will probably be the case for a while...

Instrument	Pessimistic	Moderate	Optimistic
PAPER	1.65	1.93	8.86
MWA	0.60	2.46	6.40
LOFAR NL Core	1.35	2.76	17.37
HERA	32.09	38.20	133.15
SKA1 Low Core	14.05	97.92	284.85

- the analysis requires a 12 order of magnitude (1PB -> 1kB) compression
  Examples: of things EoR is **not** like:
  - human generated data (social media)
  - simulations
  - pattern search (spam, intelligence, SETI)

- Examples of things EoR IS like:
  - Deep space communications (super extreme compression)
  - Saving for retirement

PAPER overview

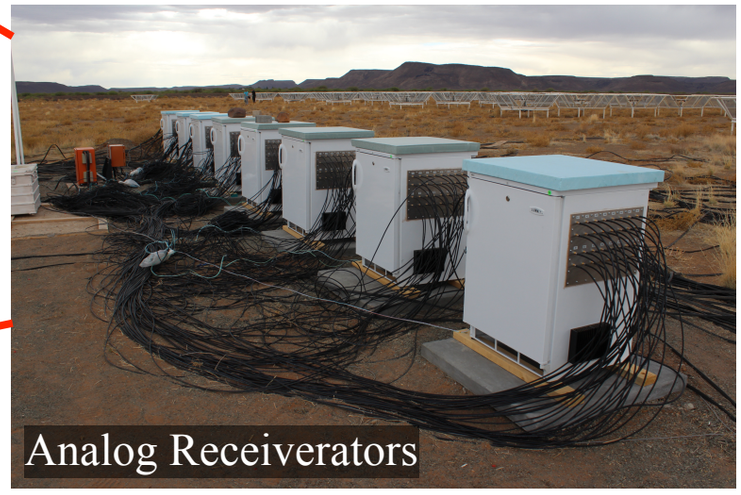

Distiller Compression Cluster

Penn Analysis Cluster

MWA Overview

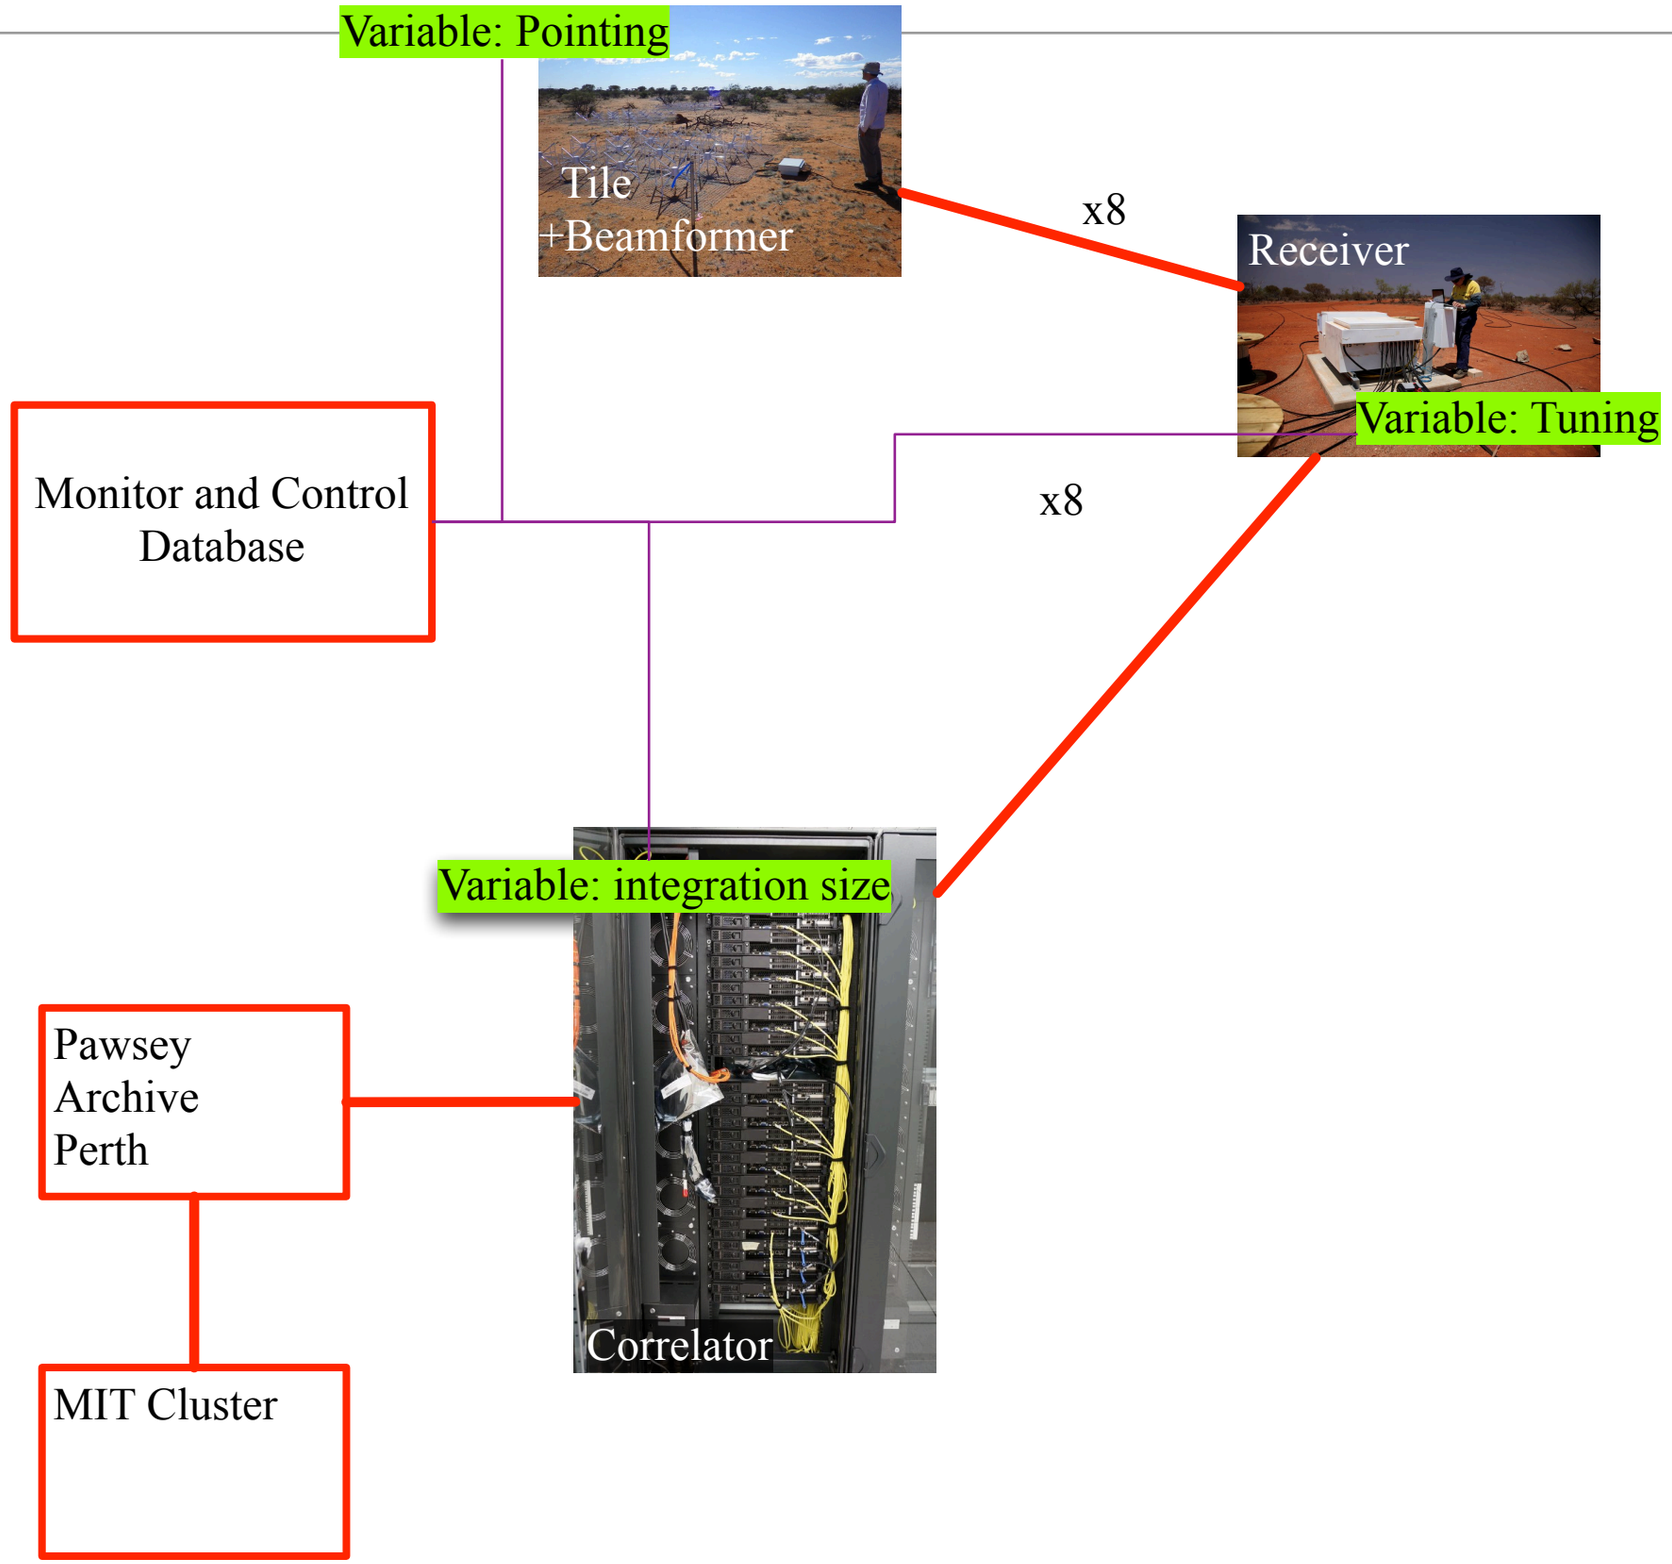

Primary MWA EoR Processing pipeline

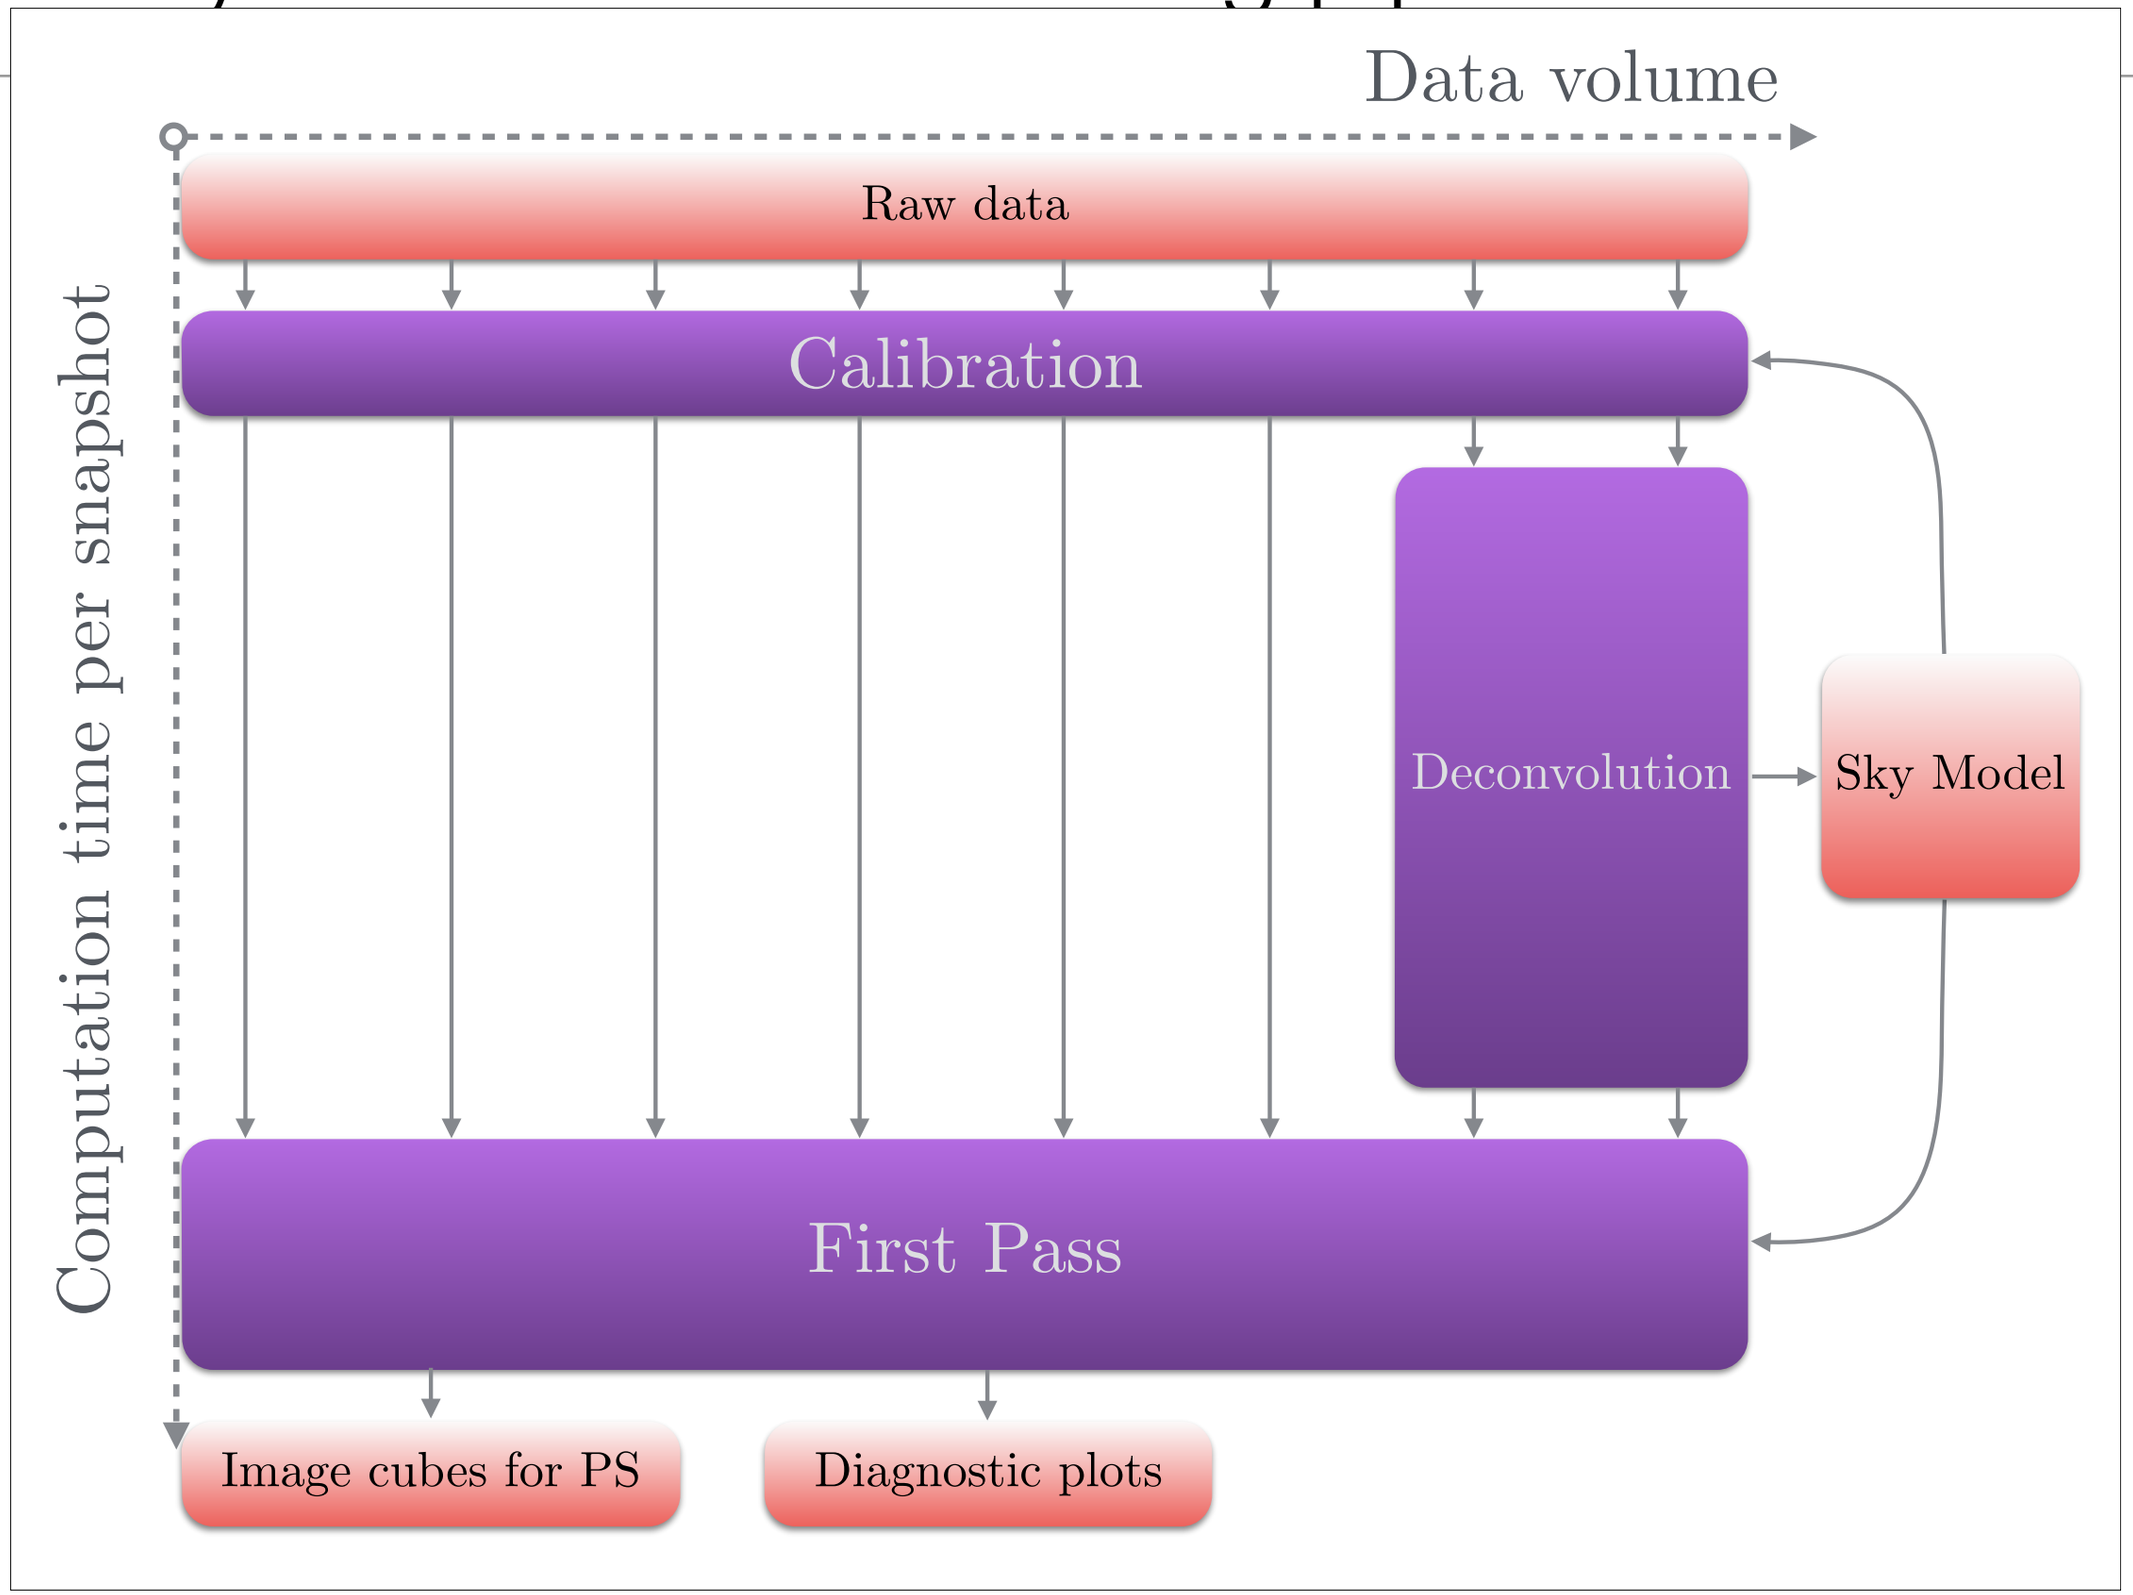

Primary PAPER Processing Pipeline

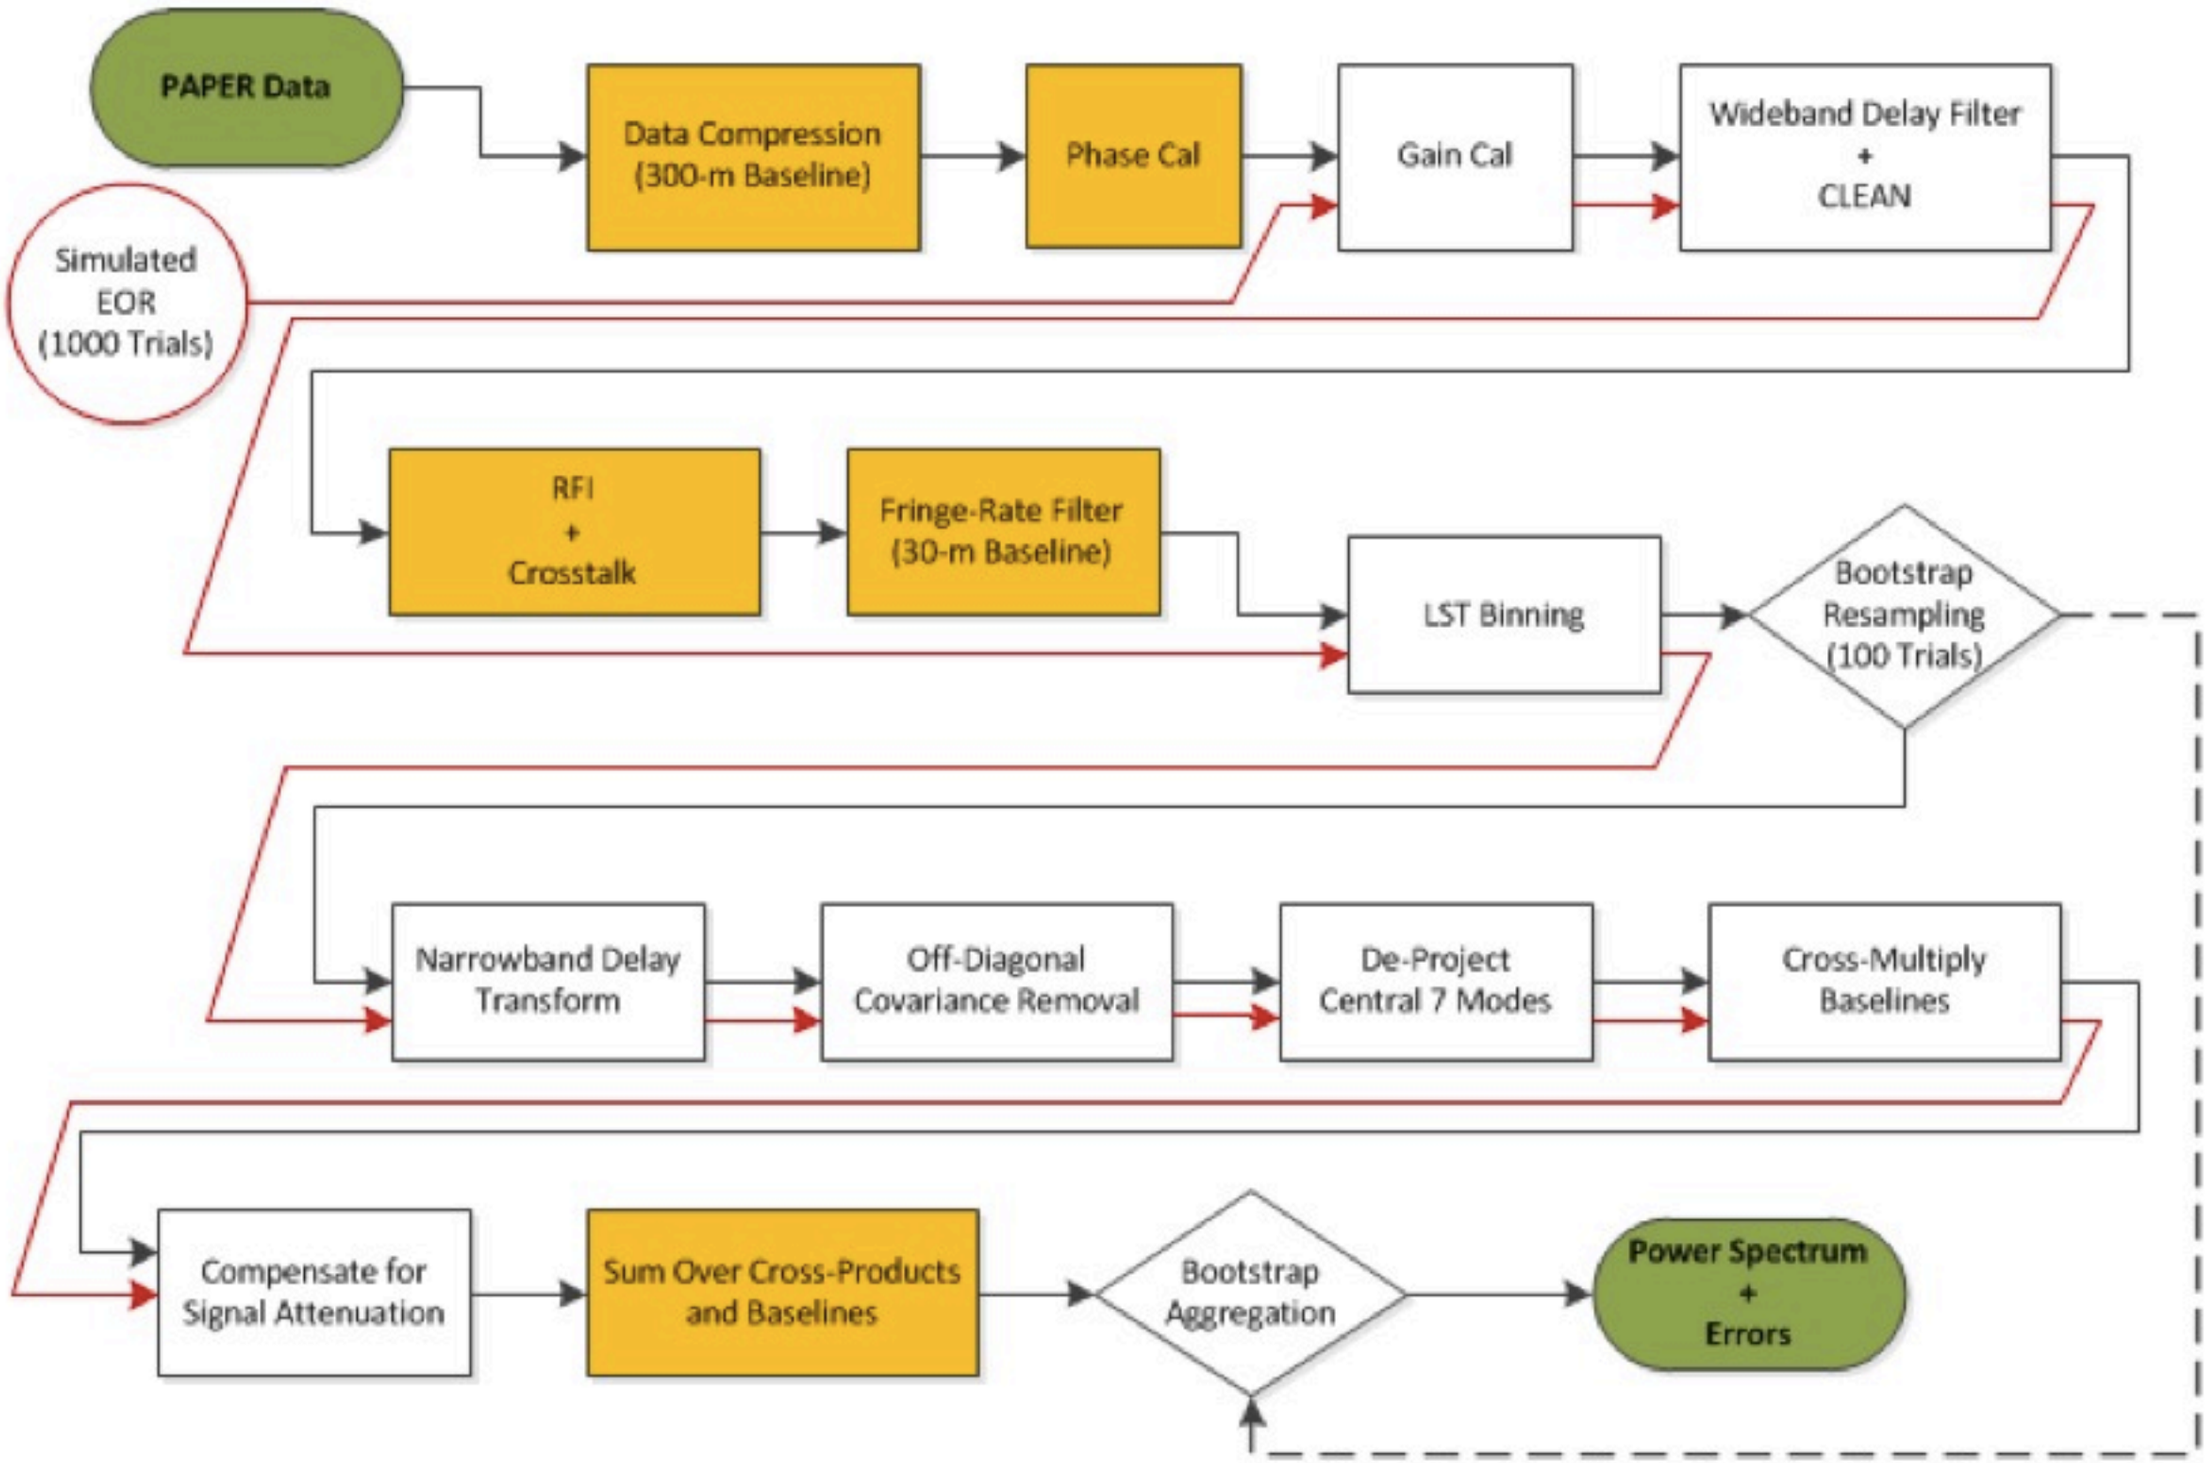

Pipeline comparison

PAPER	MWA
Raw Data (25 MBps)	Raw Data (225 MBps)
Averaging on The Distiller x70	Transport via NGAS (Internet ~300MBps)
Transport (Internet ~40MBps)	Averaging x2
Calibration	Calibration
Foreground Filter	Foreground Model Subtract
Time grid (120 seconds)	Grid (112 seconds)
LST Average (corner turn)	LST Average (bootstrap)
Power spectrum	Power spectrum

NGAS

- postgresql database of files
- web API
- scripts interact with database to:
  data replication
  archiving
  availability

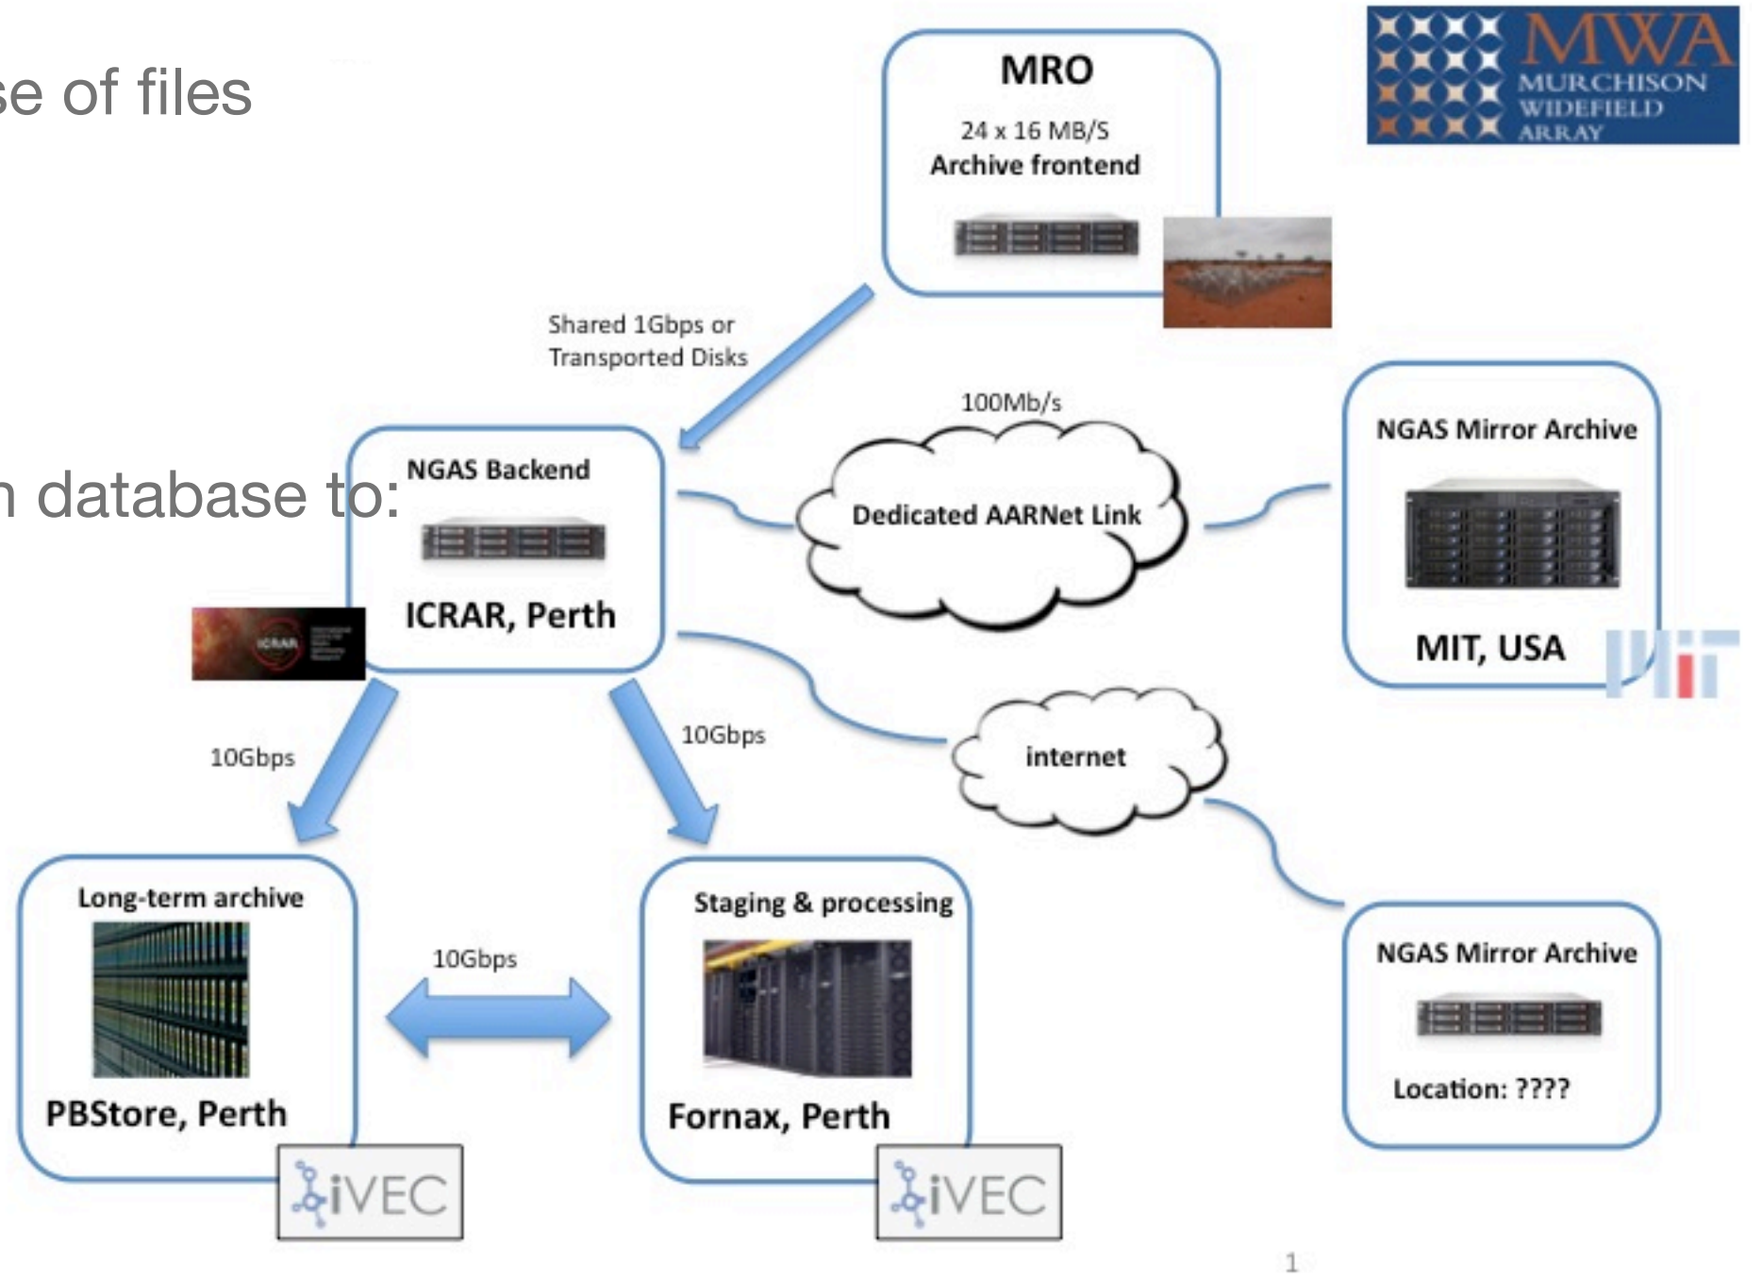

Some MWA numbers

100

74

5

14151

330

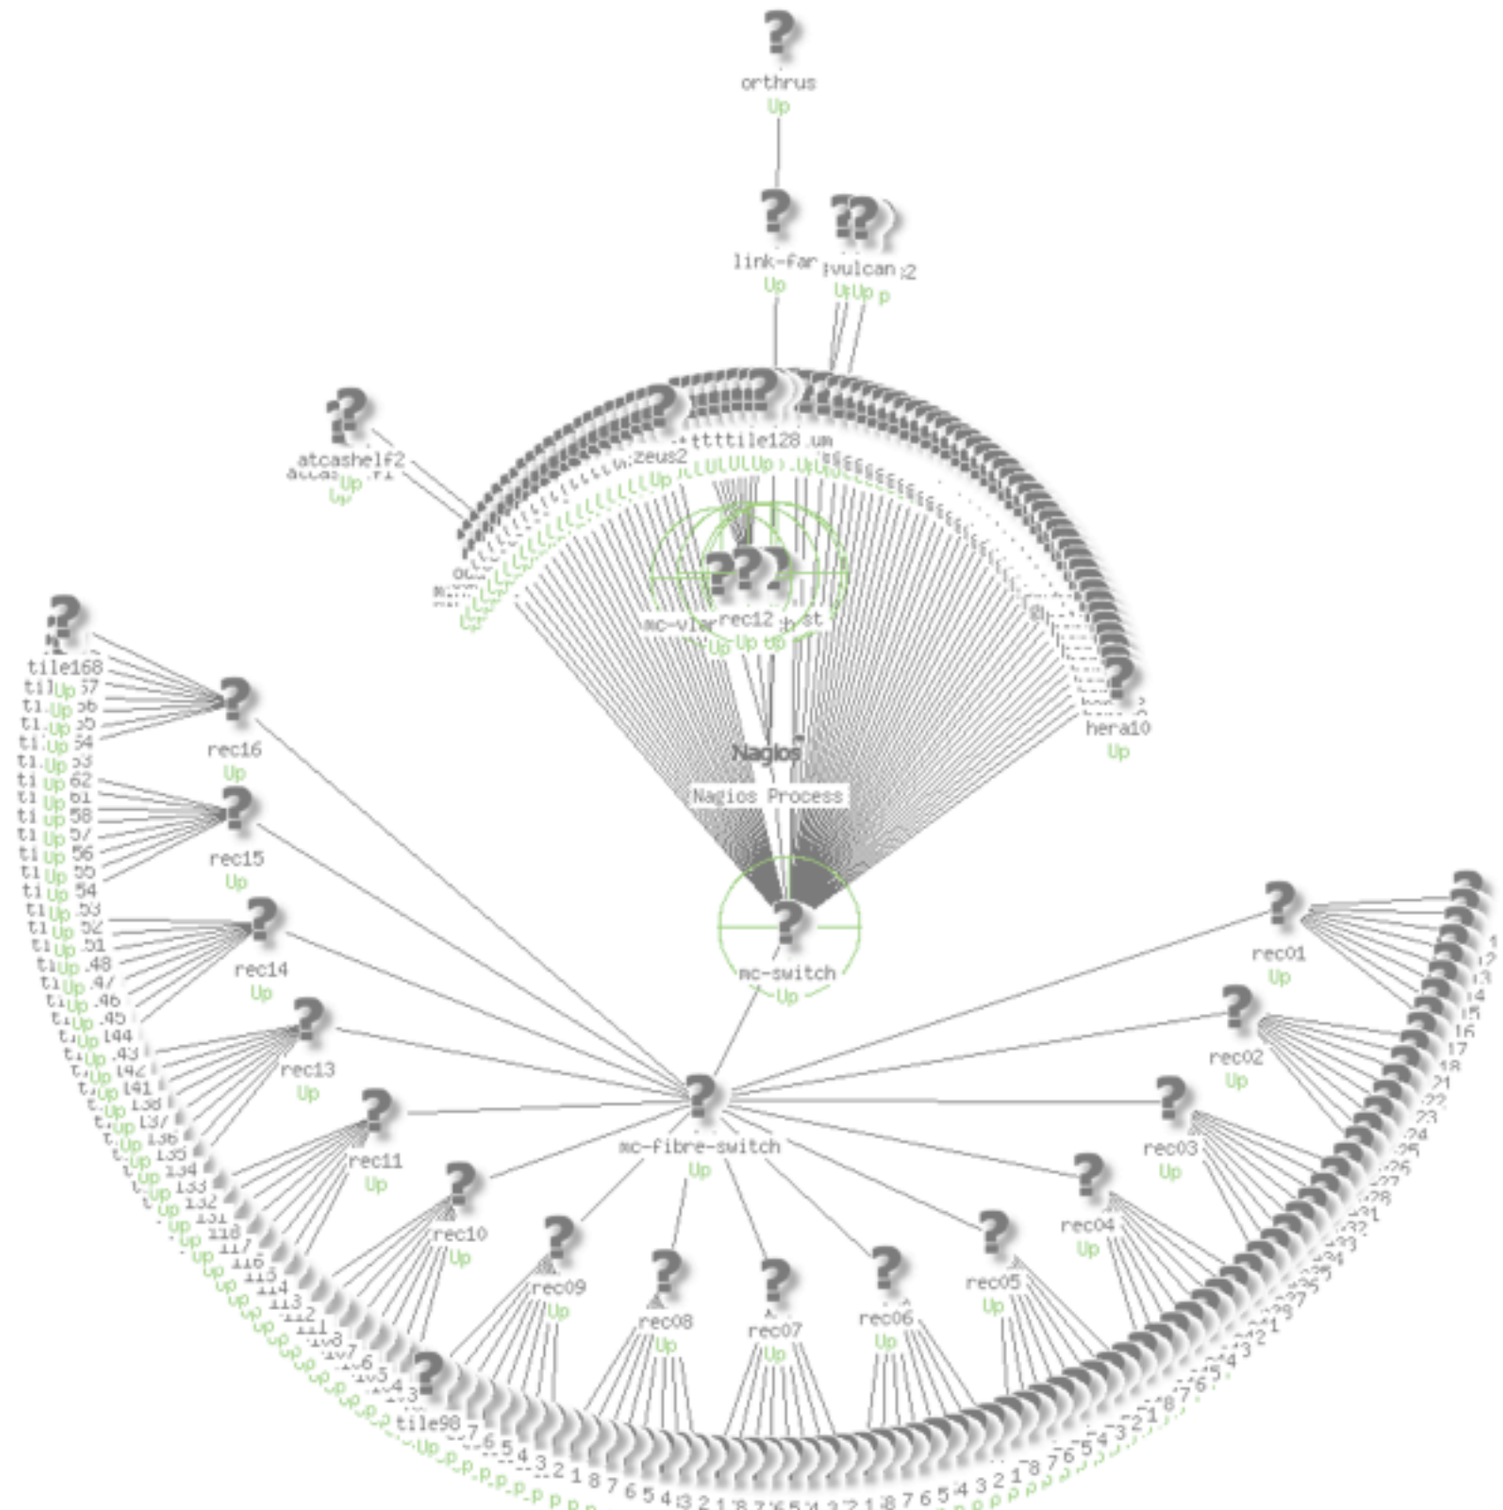

Some MWA numbers

100 independent computers in the array

74

5

14151

330

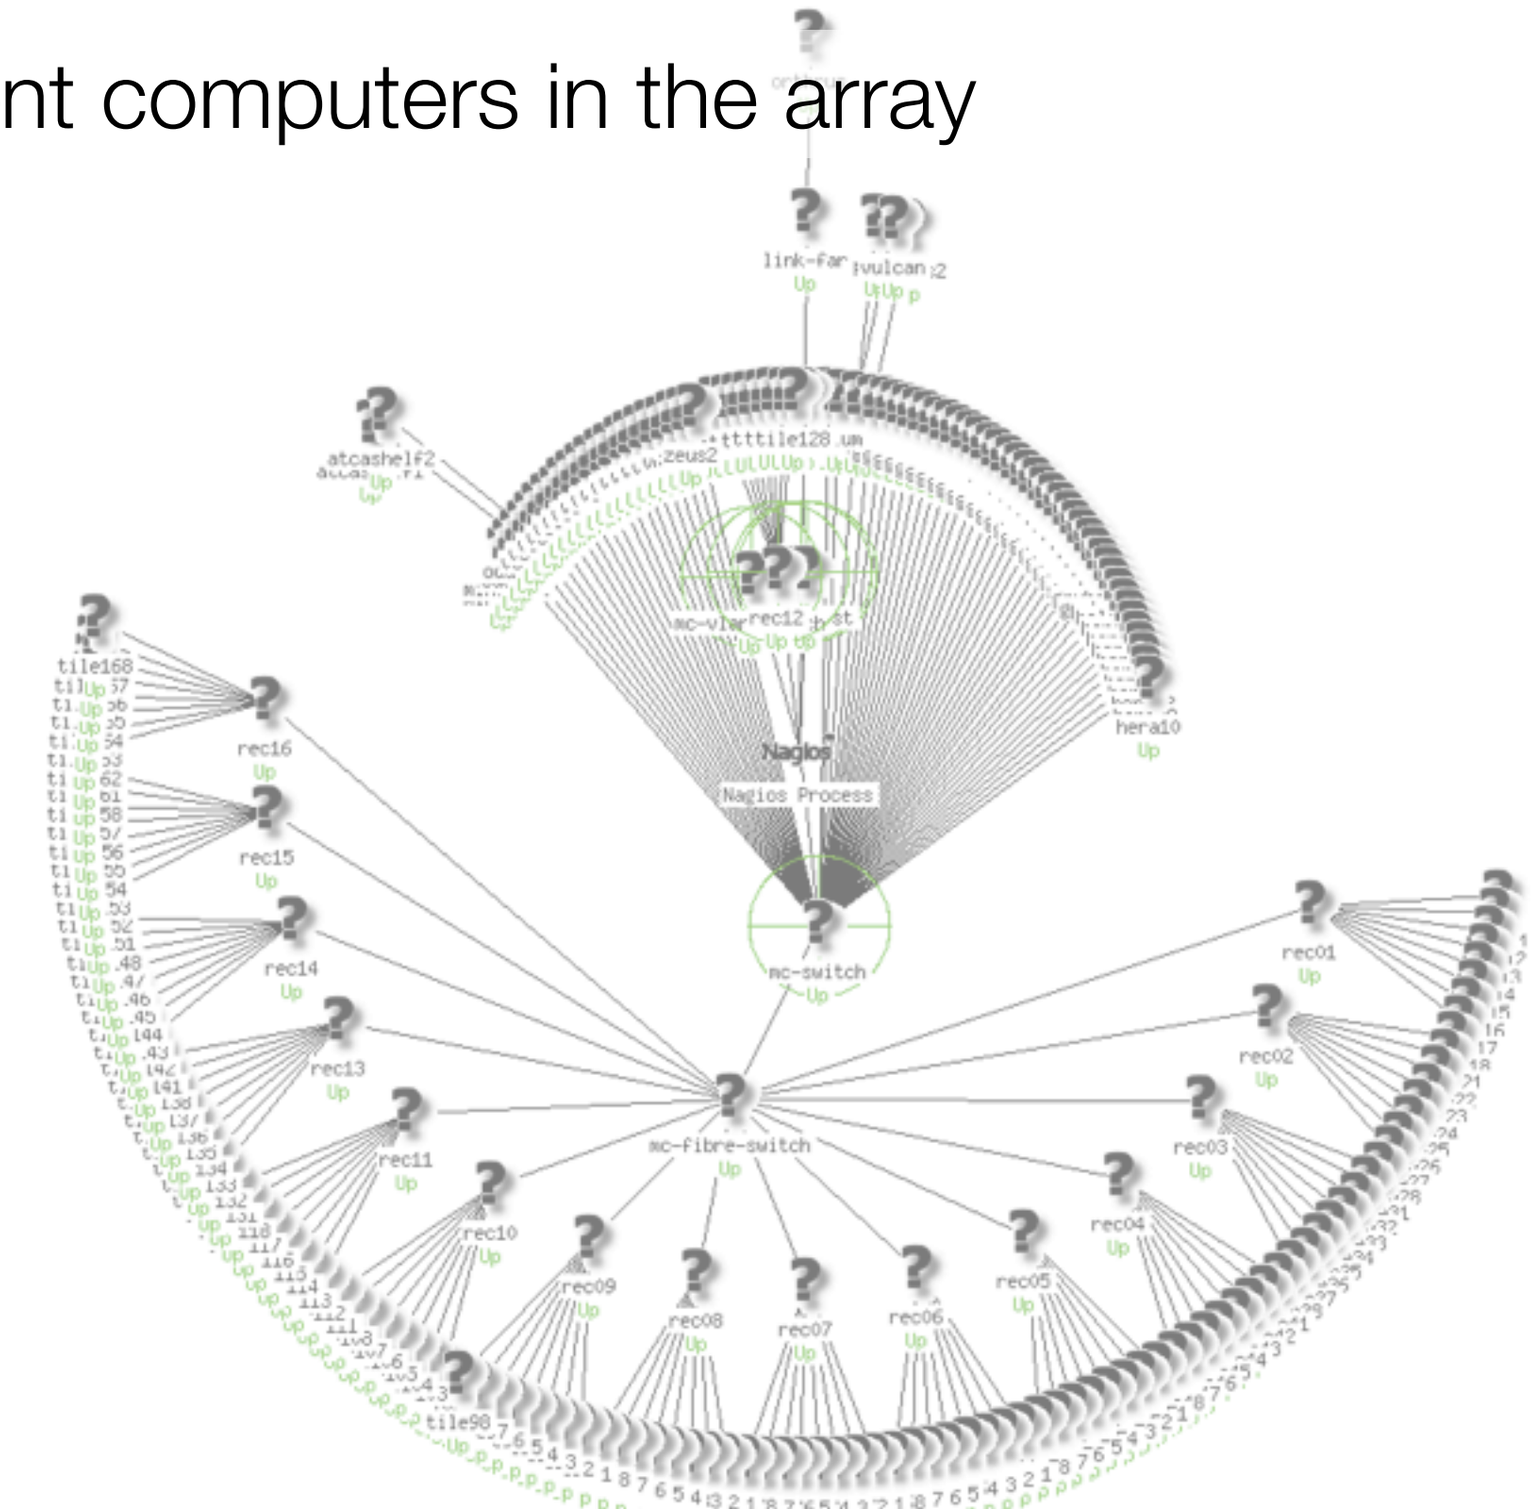

Some MWA numbers

100 independent computers in the array

74 tables in monitor and control

5

14151

330

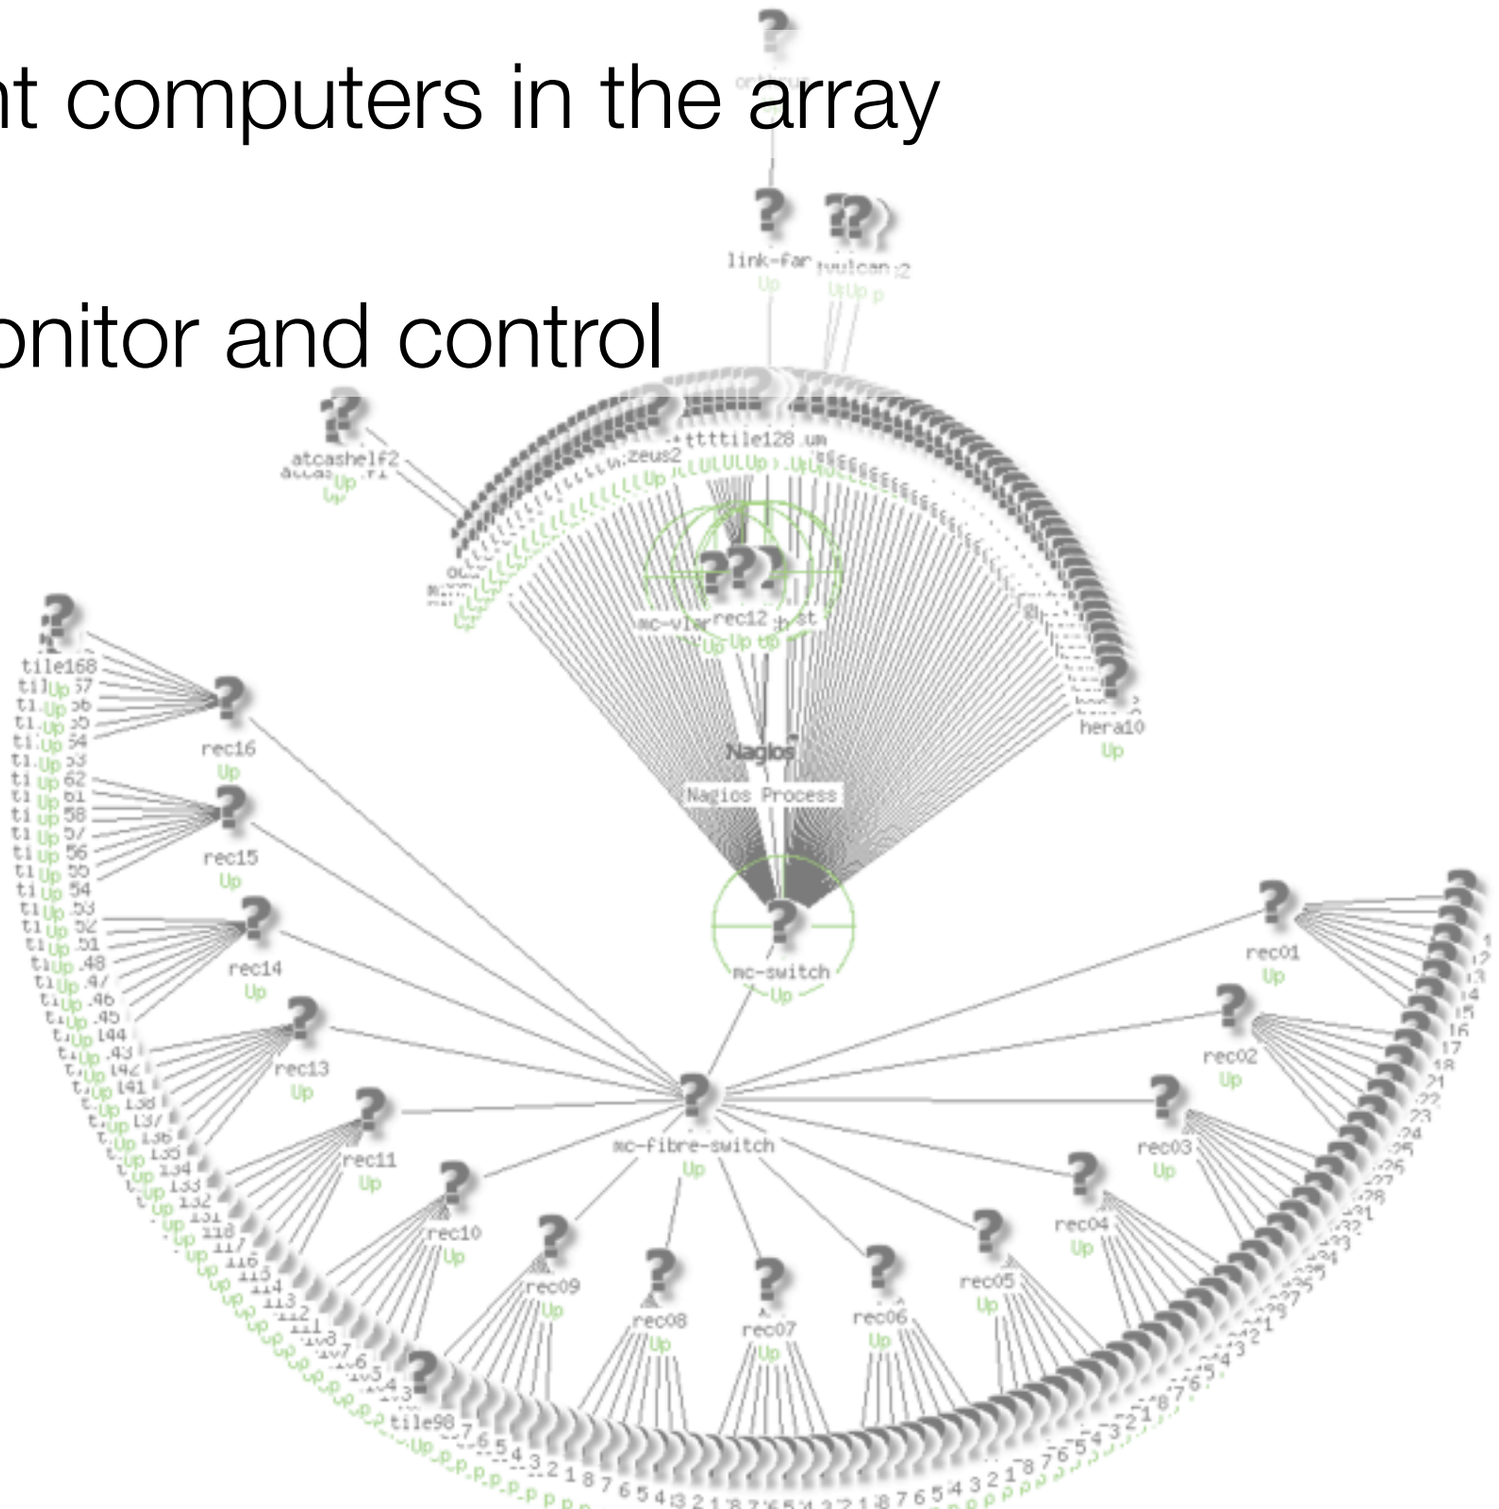

Some MWA numbers

100 independent computers in the array

74 tables in monitor and control

14151

5 key observing control parameters

330

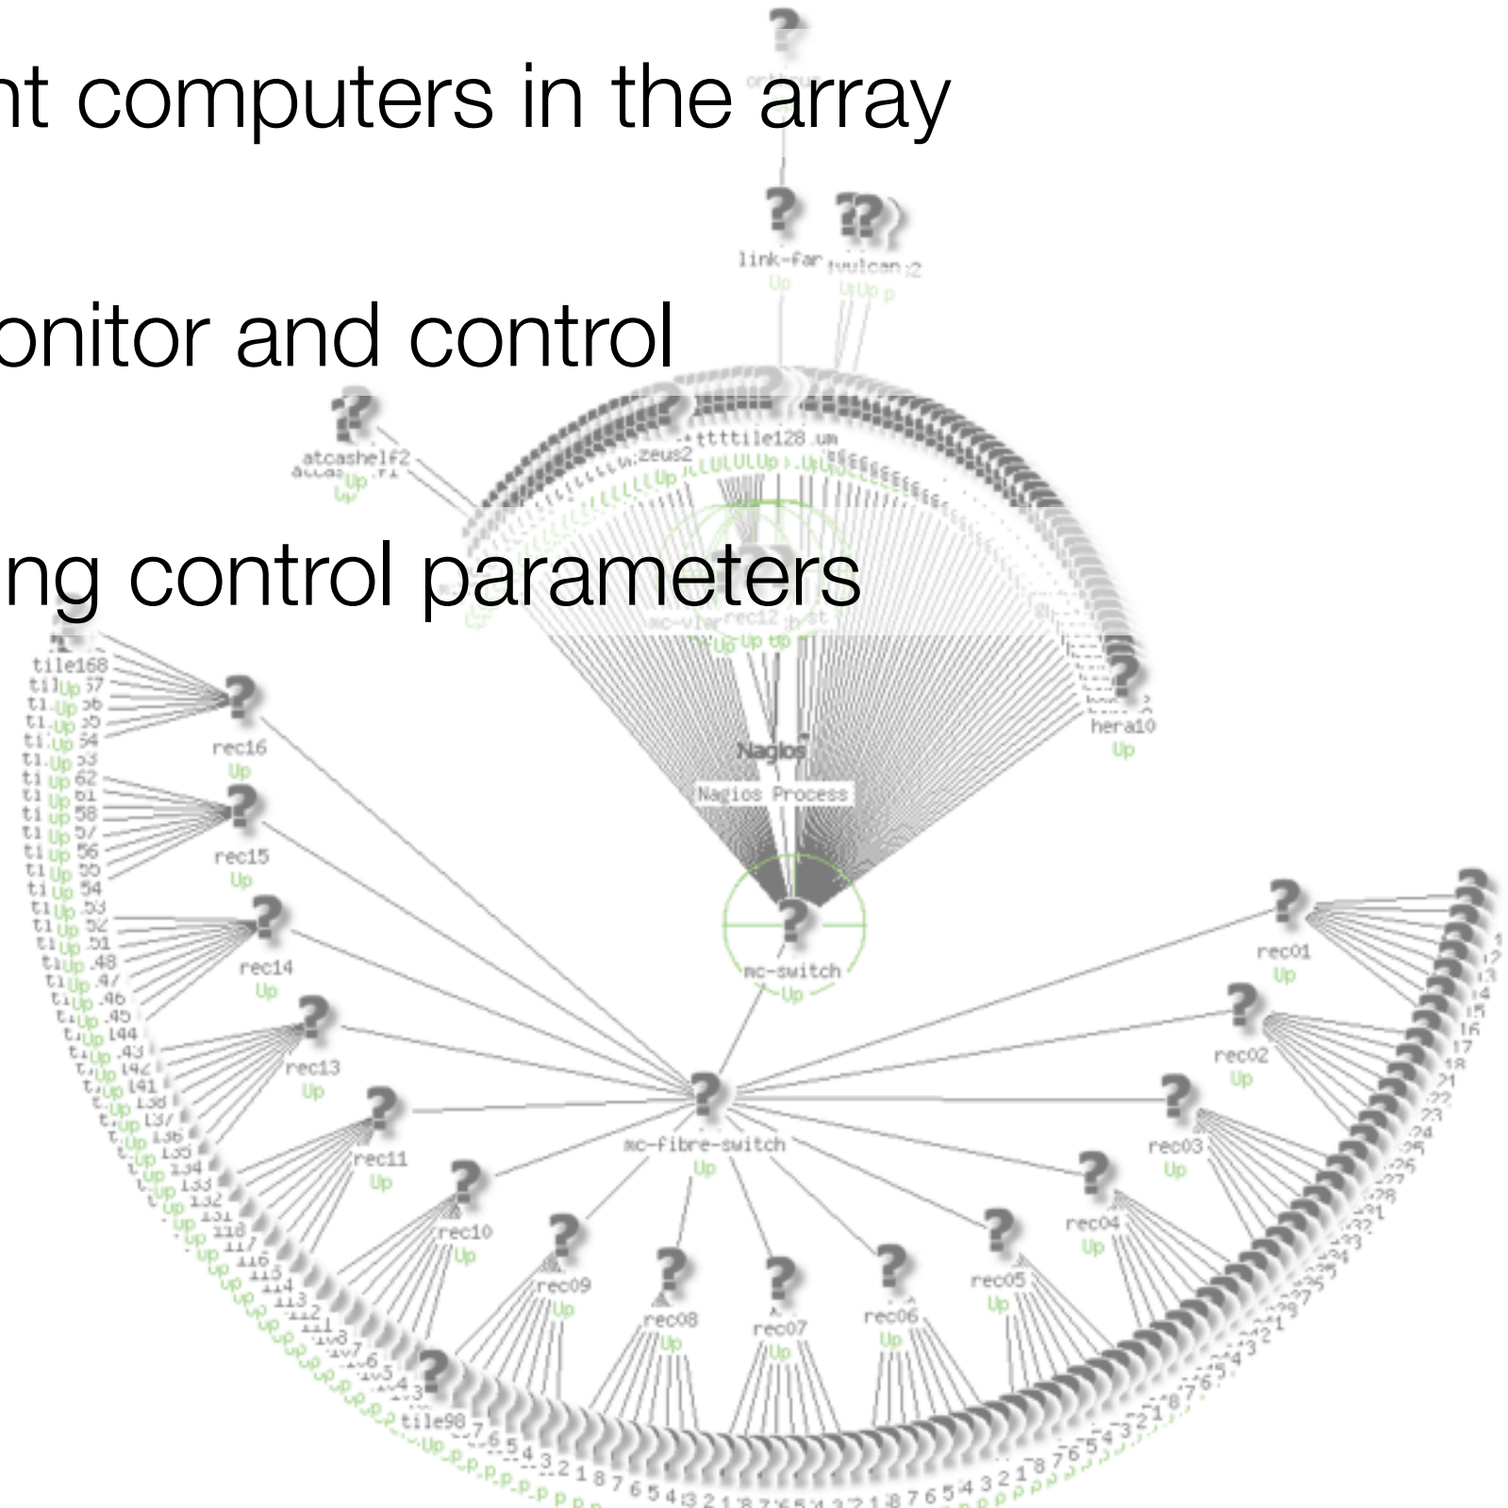

Some MWA numbers

100 independent computers in the array

74 tables in monitor and control

5 key observing control parameters

14151 observations since July 1, 2013

330

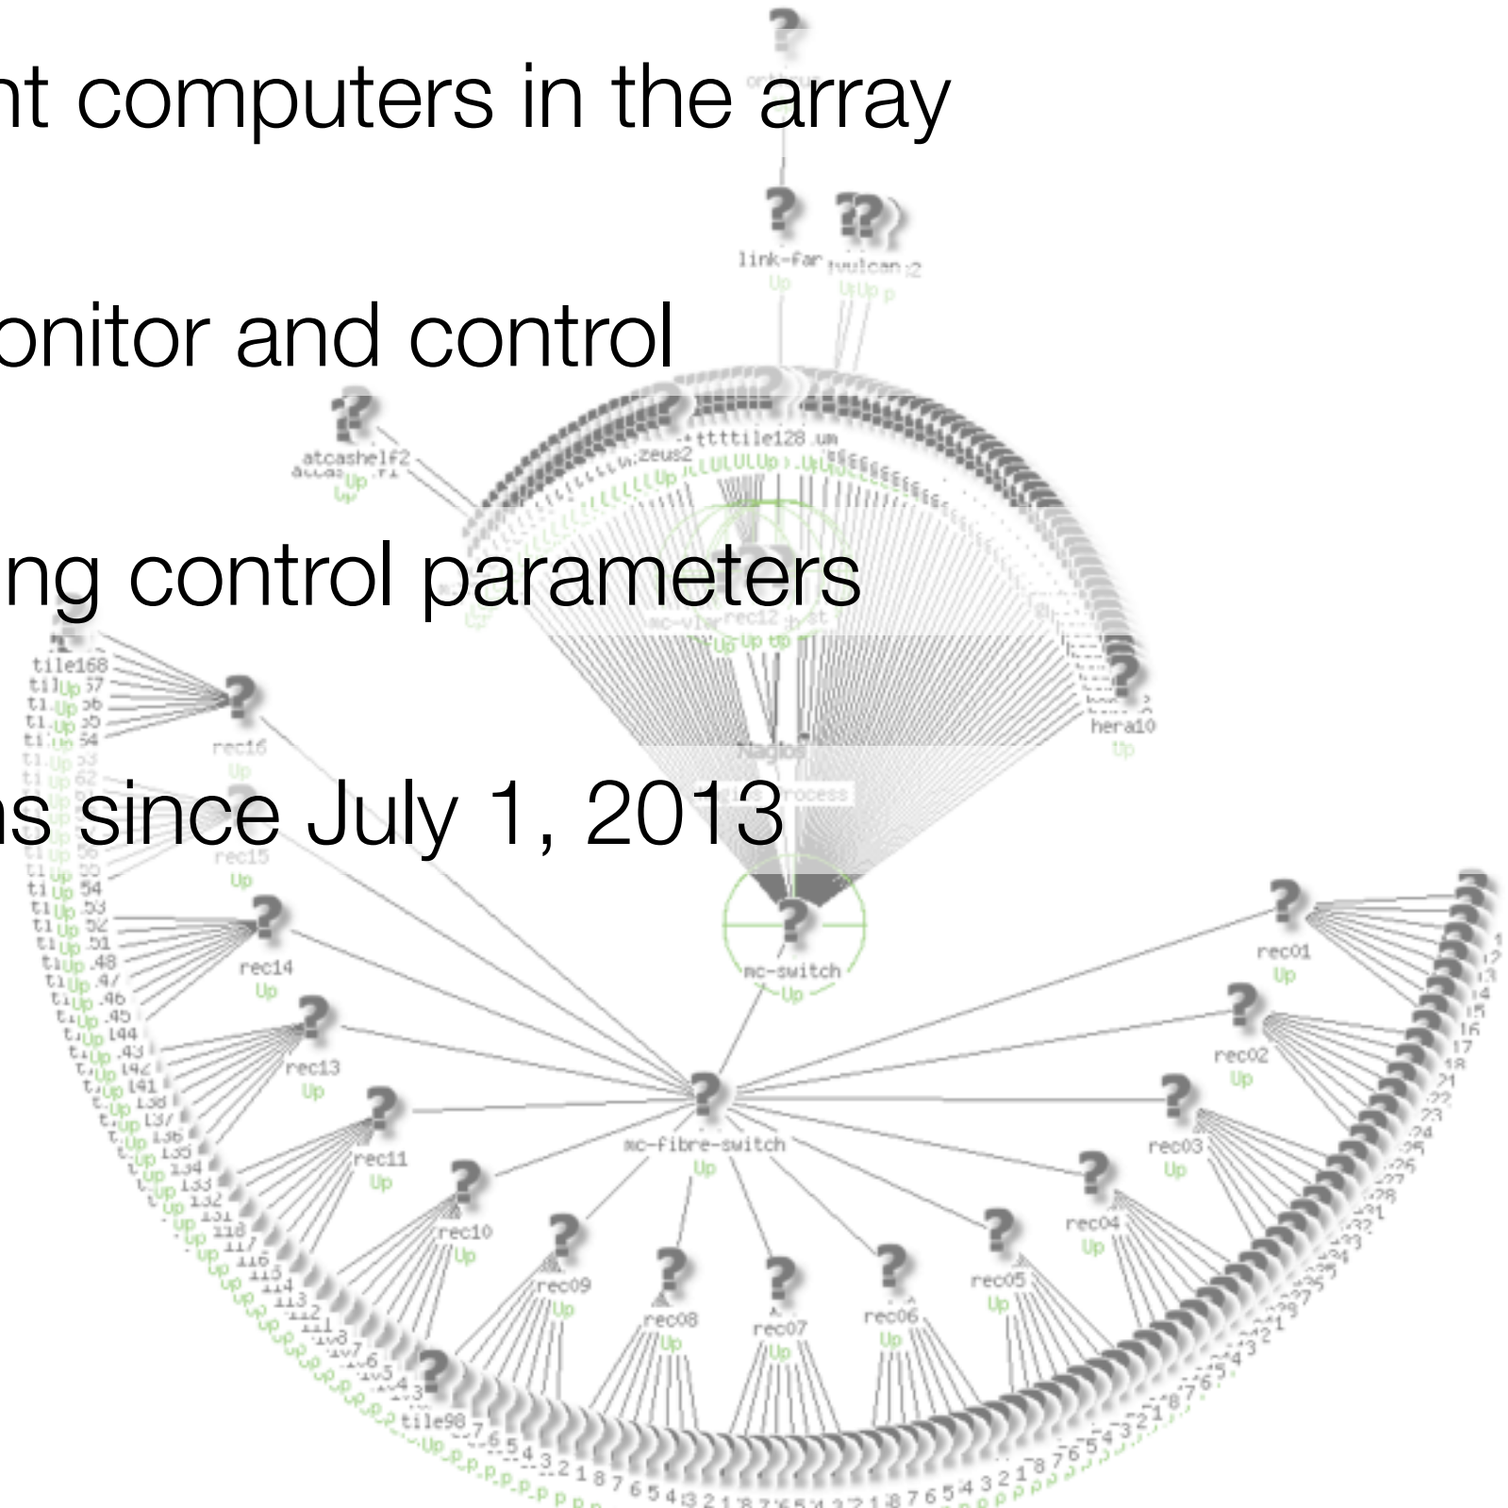

Some MWA numbers

100 independent computers in the array

74 tables in monitor and control

5 key observing control parameters

14151 observations since July 1, 2013

330 TB of EoR per person

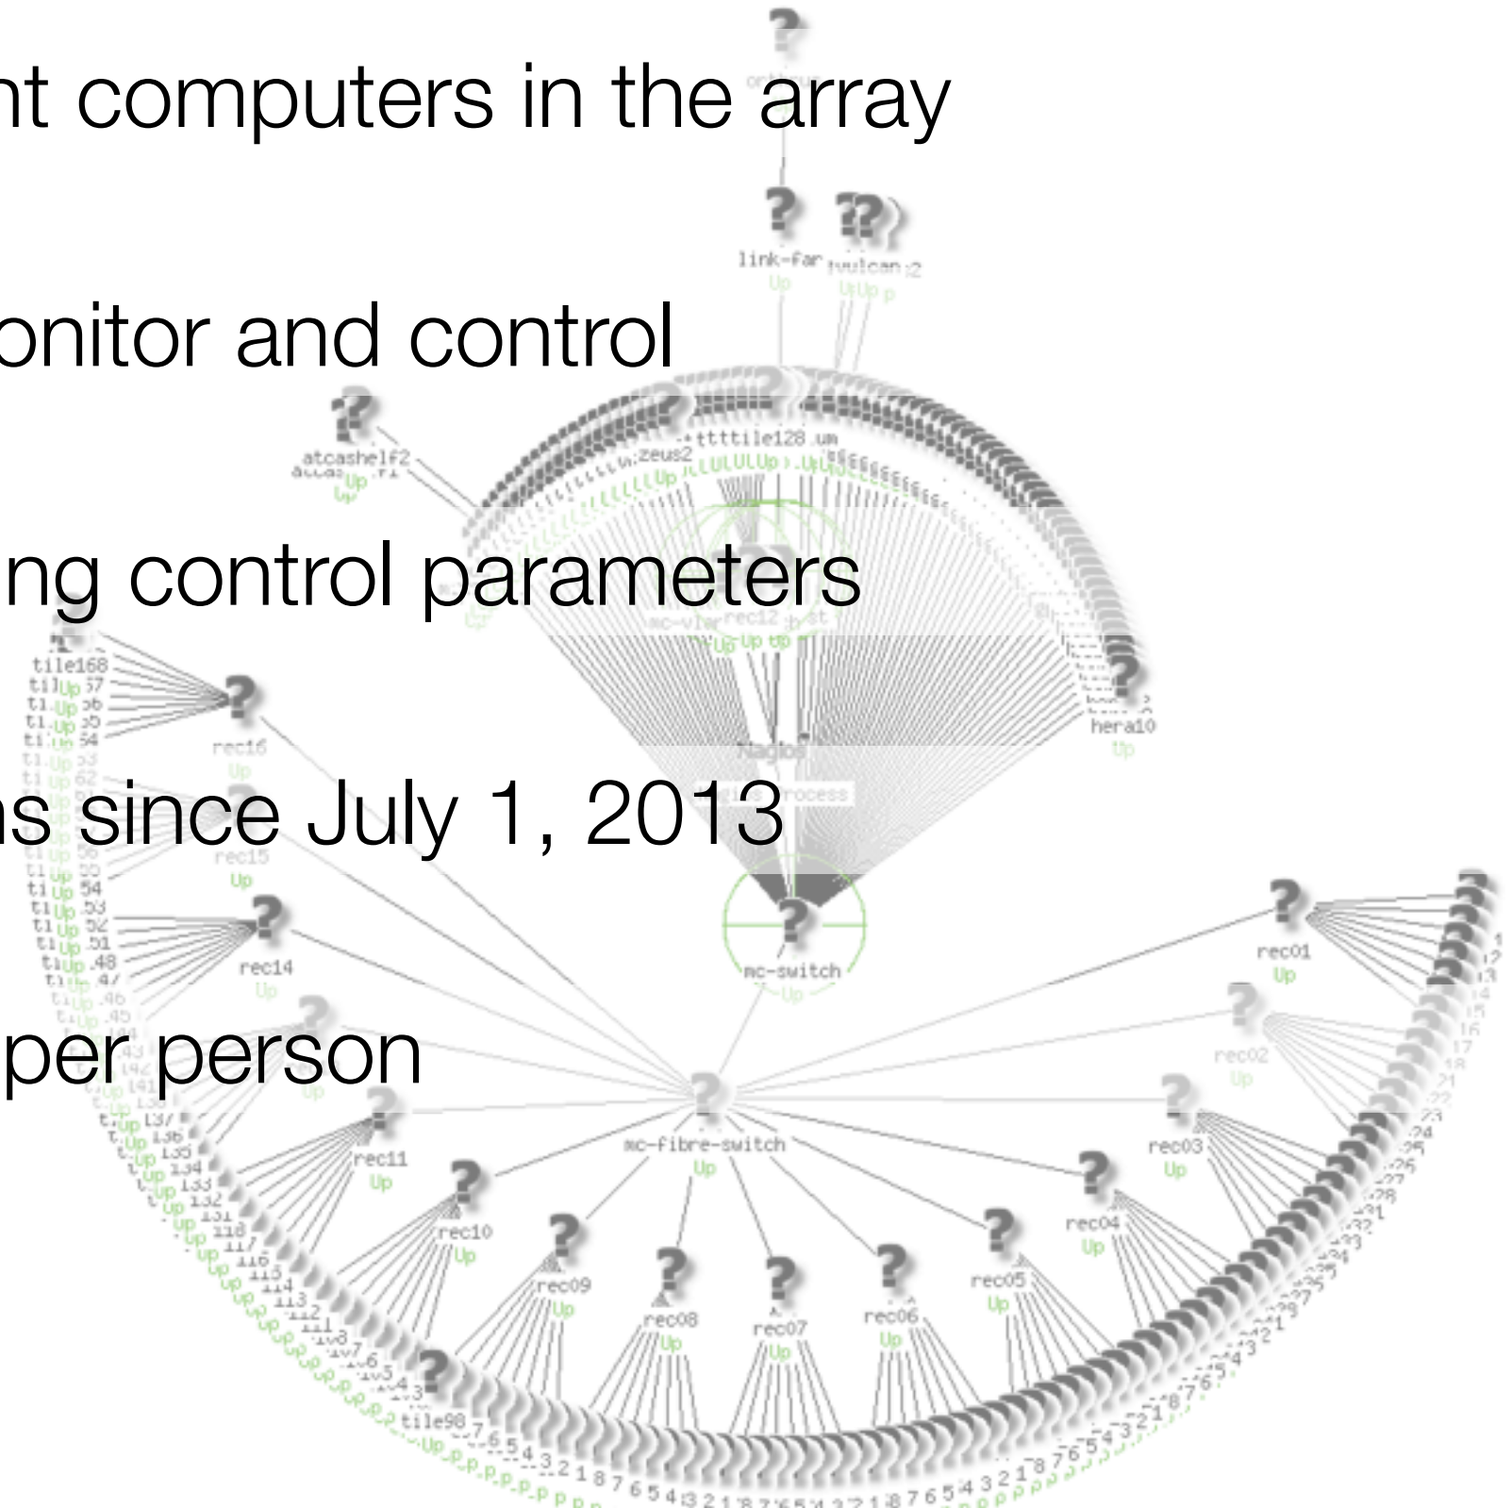

Some PAPER numbers

7

0.75

15000

6

0.183$

3

Some PAPER numbers

7	Number of new arrays commissioned
0.75	
15000	
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	Average number of times data moved during operations
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	Average number of times data moved during operations
0.183$	Average price paid per GB of storage
3	

ORACLE

Sun ORACLE

Complete storage. Better information.

Tiered Storage Savings Calculator

Overview | Analysis | Next Steps | Assumptions

Step 1: Enter information about your company:

Please enter total terabytes created per year. (i) 1000

Please enter your current average cost per gigabyte. (i) 0.138

Step 2: See your savings. Calculate

Total Savings Graph

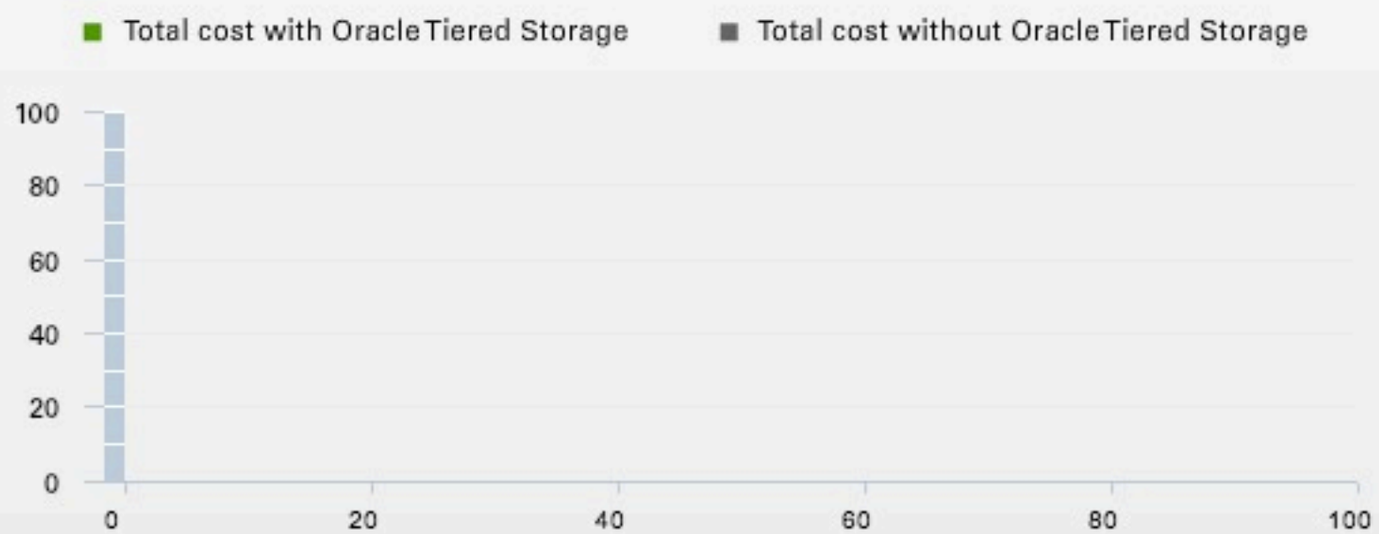

Improve Your Storage

Oracle 1-800-633-0738

- Have Oracle call you
- Global contacts
- Sales Chat Live

Tiered Storage Savings

Select Year: 7-Year Total Savings

Total Savings

Graph Savings

Hardware Savings

Graph Savings

Power Savings

Graph Savings

ORACLE

Sun ORACLE

Complete storage. Better information.

Tiered Storage Savings Calculator

Overview | Analysis | Next Steps | Assumptions

Step 1: Enter information about your company:

Please enter total terabytes created per year. 1000

Please enter your current average cost per gigabyte. 3.6

Step 2: See your savings. Calculate

Total Savings Graph

Total cost with Oracle Tiered Storage

Total cost without

36
30
24
18
12
6
0

2010 2011 2012 2013 2014 2015 2016

Please enter a number above 3.5

OK

Improve Your Storage

Oracle 1-800-633-0738

- Have Oracle call you
- Global contacts
- Sales Chat Live

Tiered Storage Savings

Select Year: 7-Year Total Savings

Total Savings

$31,826,649 Graph Savings

Hardware Savings

$33,184,760 Graph Savings

Power Savings

$165,889 Graph Savings

Some PAPER numbers

7

0.75

15000

6

0.183$

3

Some PAPER numbers

7	Number of new arrays commissioned
0.75	
15000	
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	Average number of times data moved during operations
0.183$	
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	Average number of times data moved during operations
0.183$	Average price paid per GB of storage
3	

Some PAPER numbers

7	Number of new arrays commissioned
0.75	Average number of FTEs doing commissioning
15000	Approximate number of hours recorded to date
6	Average number of times data moved during operations
0.183$	Average price paid per GB of storage
3	number of data loss “events”

Typical Problem Areas

- transitioning to HPC with lots of data
- situational awareness

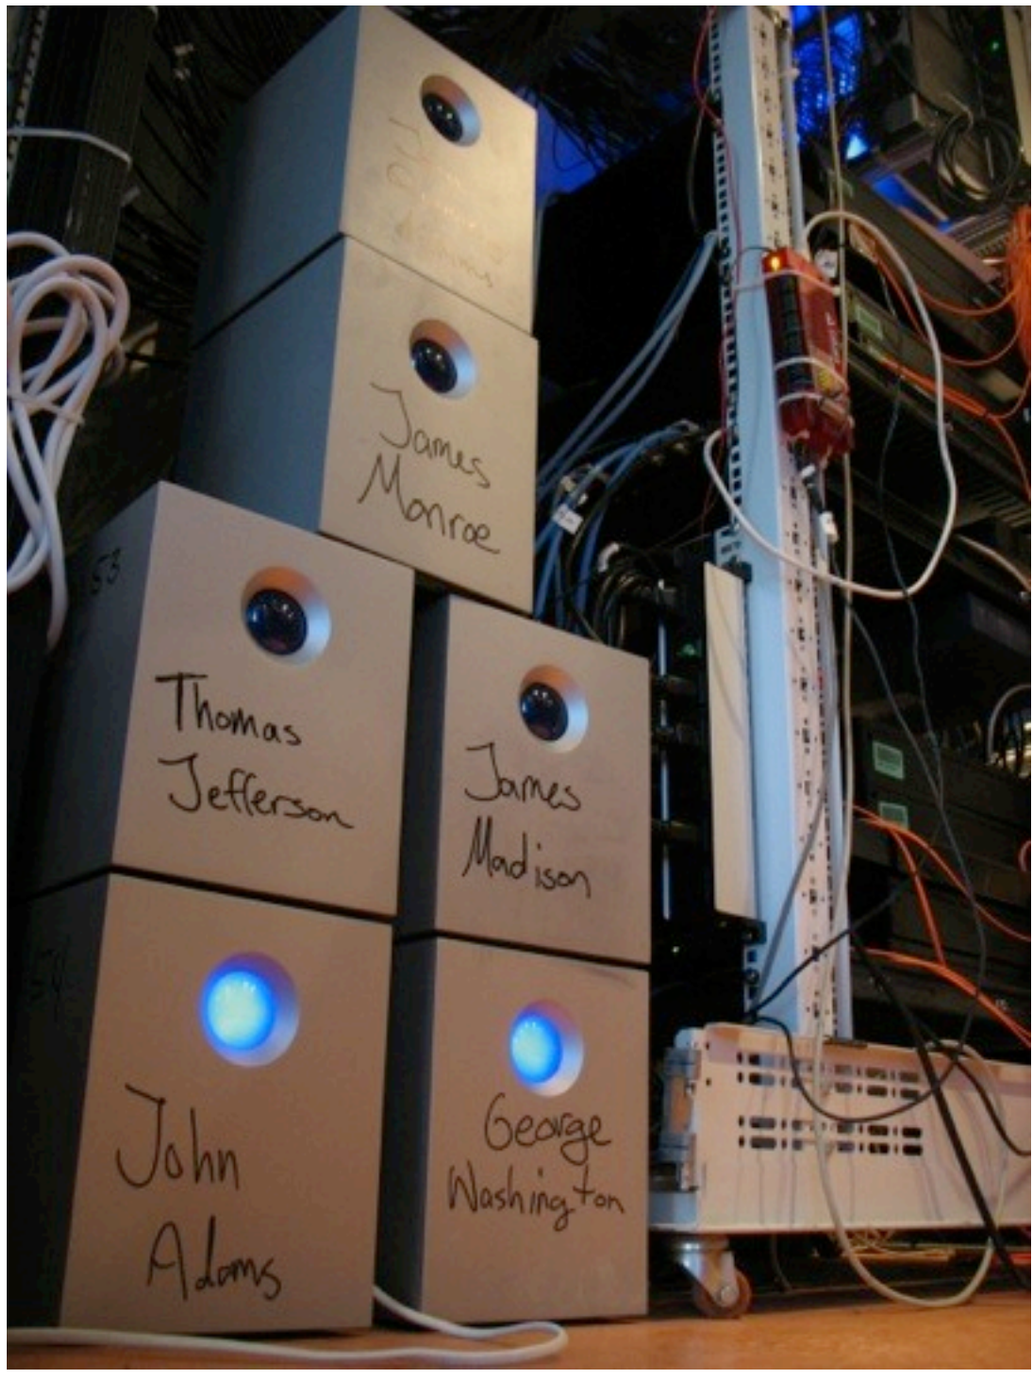

Situational Awareness

- Operations
  - Scheduling snafus
  - firewalls
  - turnaround time
  - whut is my telescope doing?

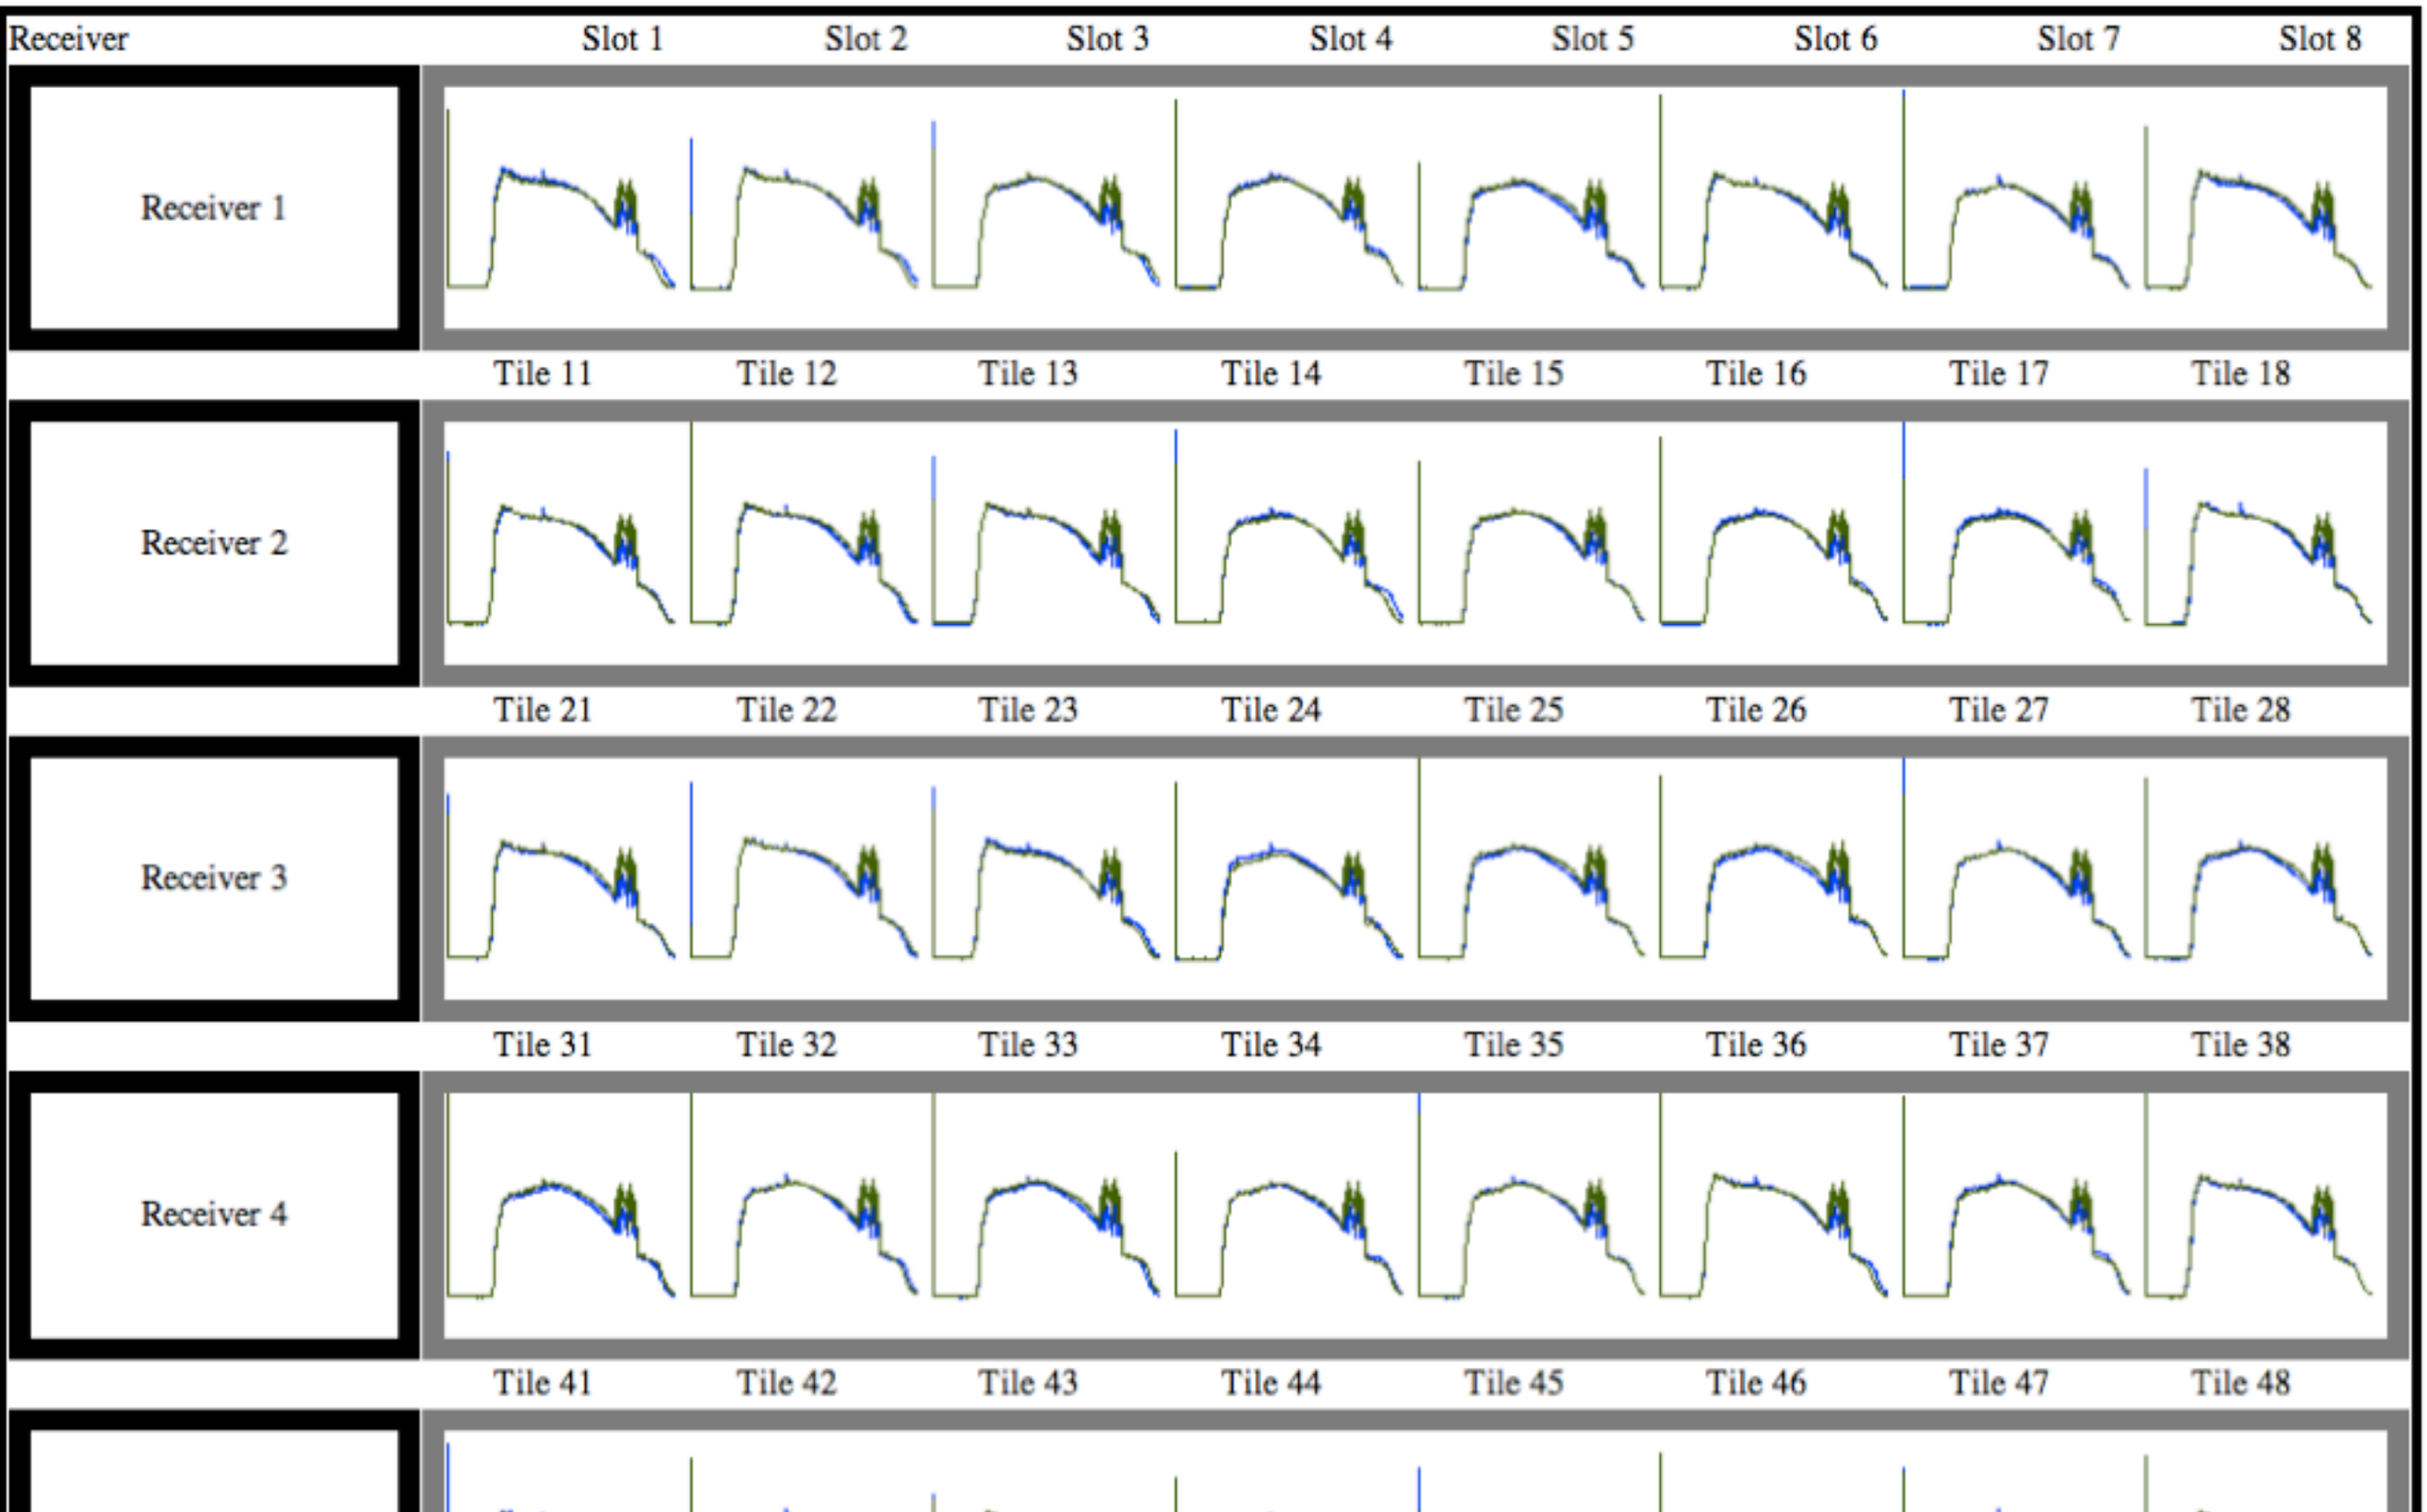
Receiver
Slot 1
Slot 2
Slot 3
Slot 4
Slot 5
Slot 6
Slot 7
Slot 8
Receiver 1
Tile 11
Tile 12
Tile 13
Tile 14
Tile 15
Tile 16
Tile 17
Tile 18
Receiver 2
Tile 21
Tile 22
Tile 23
Tile 24
Tile 25
Tile 26
Tile 27
Tile 28
Receiver 3
Tile 31
Tile 32
Tile 33
Tile 34
Tile 35
Tile 36
Tile 37
Tile 38
Receiver 4
Tile 41
Tile 42
Tile 43
Tile 44
Tile 45
Tile 46
Tile 47
Tile 48

Receiver 13

Tile 131 Tile 132 Tile 133 Tile 134 Tile 135 Tile 136 Tile 137 Tile 138

Receiver 14

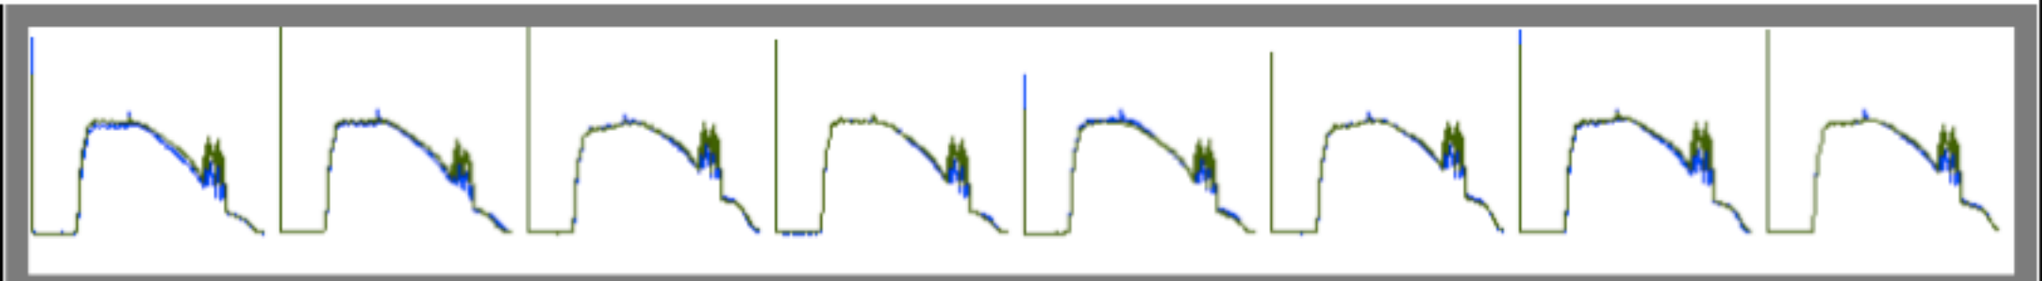

Tile 141 Tile 142 Tile 143 Tile 144 Tile 145 Tile 146 Tile 147 Tile 148

Receiver 13

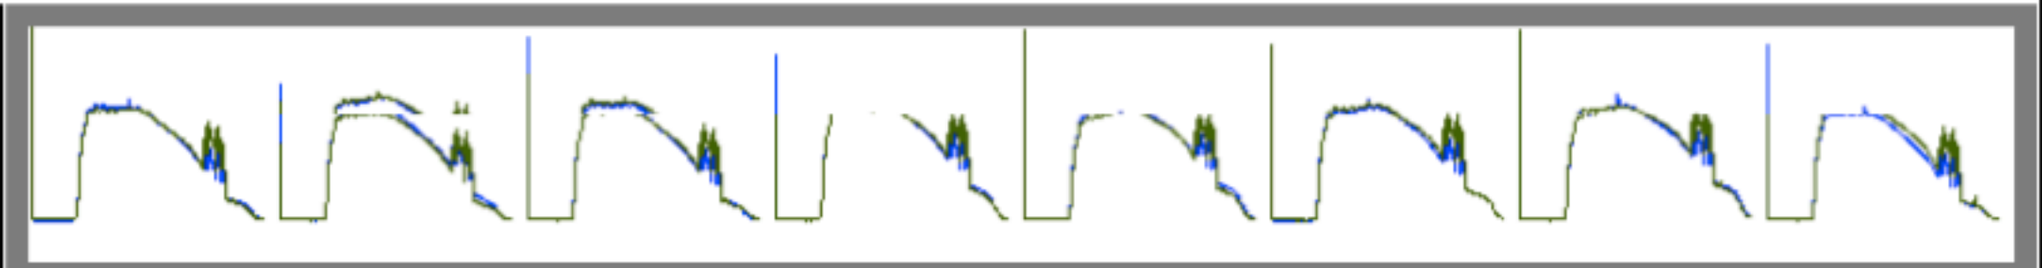

Tile 131 Tile 132 Tile 133 Tile 134 Tile 135 Tile 136 Tile 137 Tile 138

Receiver 14

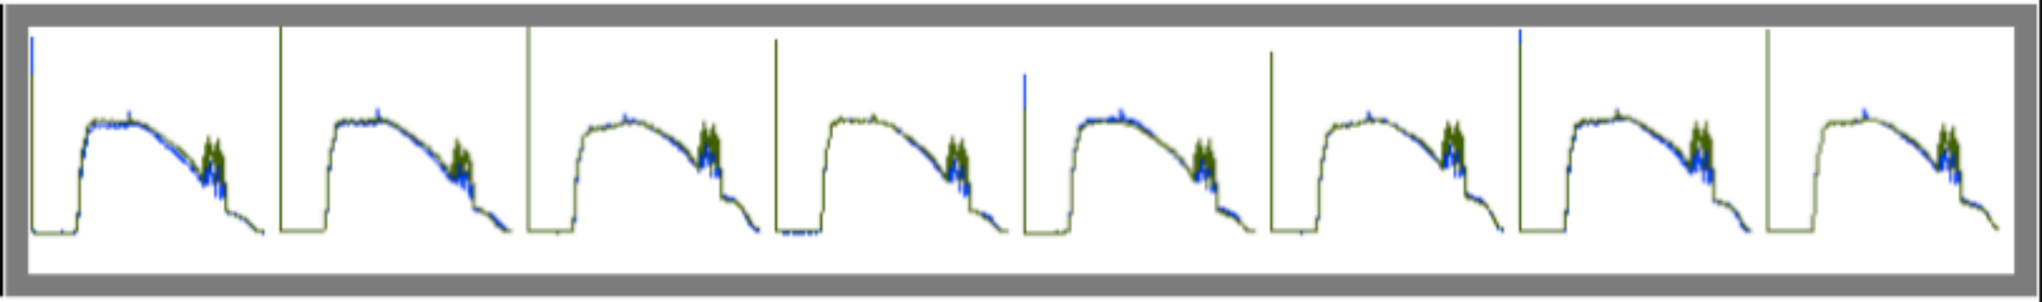

Tile 141 Tile 142 Tile 143 Tile 144 Tile 145 Tile 146 Tile 147 Tile 148

Receiver 15

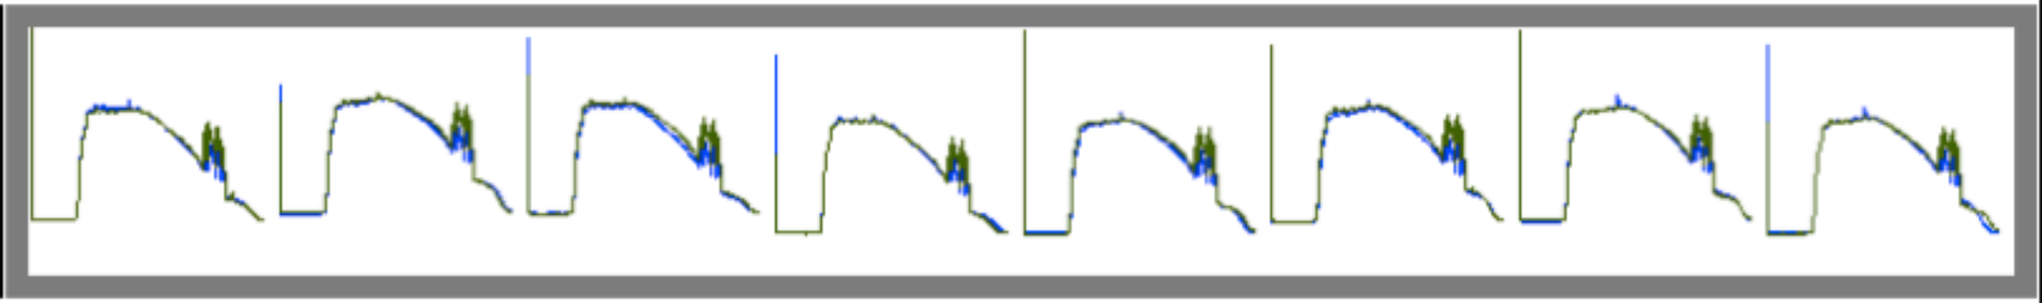

Tile 151 Tile 152 Tile 153 Tile 154 Tile 155 Tile 156 Tile 157 Tile 158

Receiver 16

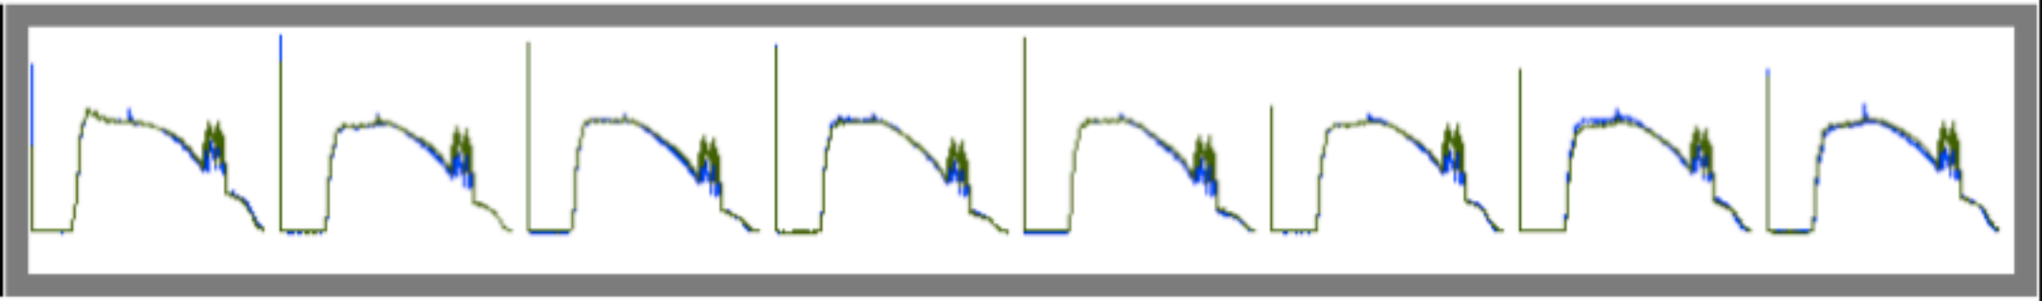

Tile 161 Tile 162 Tile 163 Tile 164 Tile 165 Tile 166 Tile 167 Tile 168

Webpage Screenshot

PFB power plots for the 12 hours ending at Wed Mar 26 18:00:01 2014 AWST

Rec01

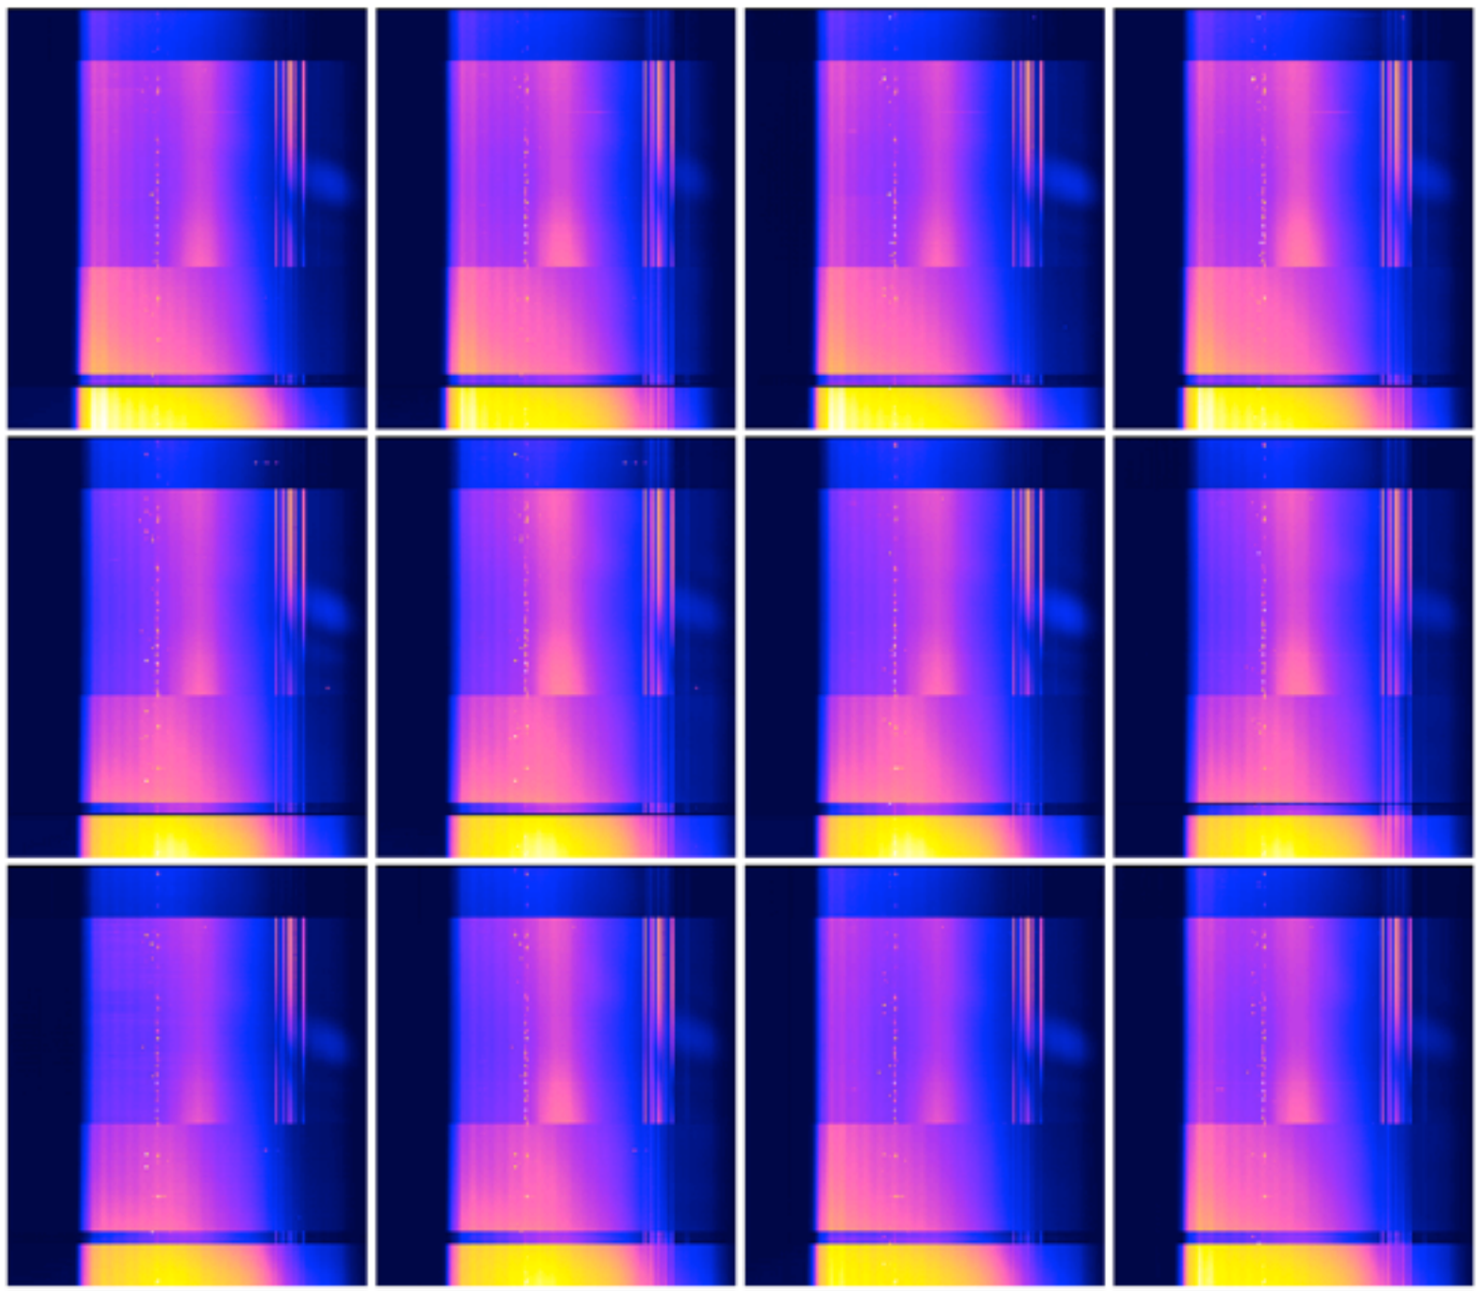

Rec16

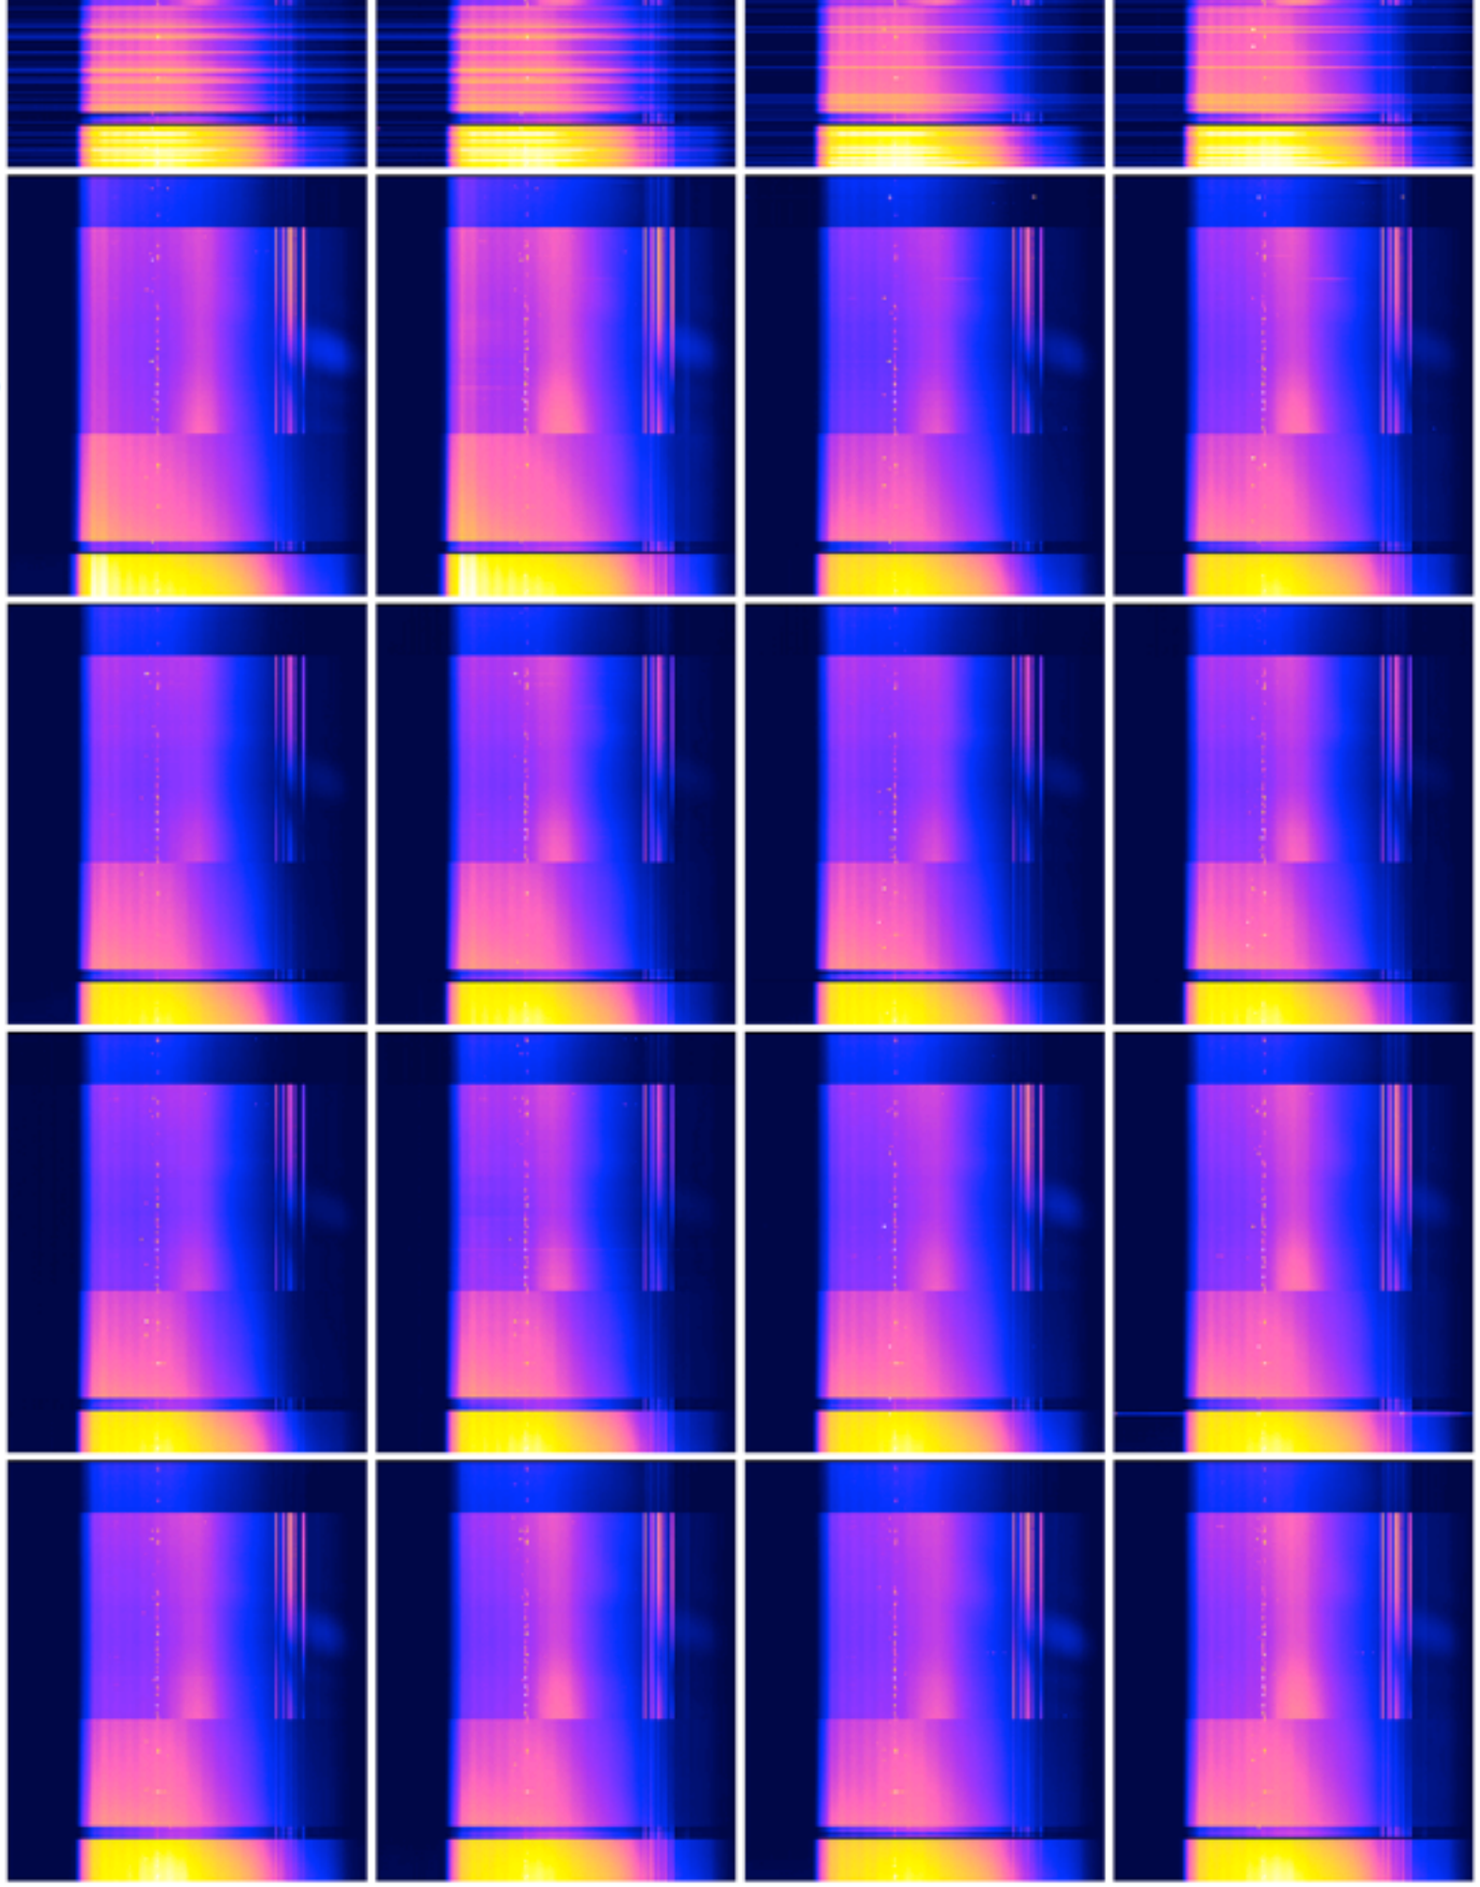

Trying to mine the metadata

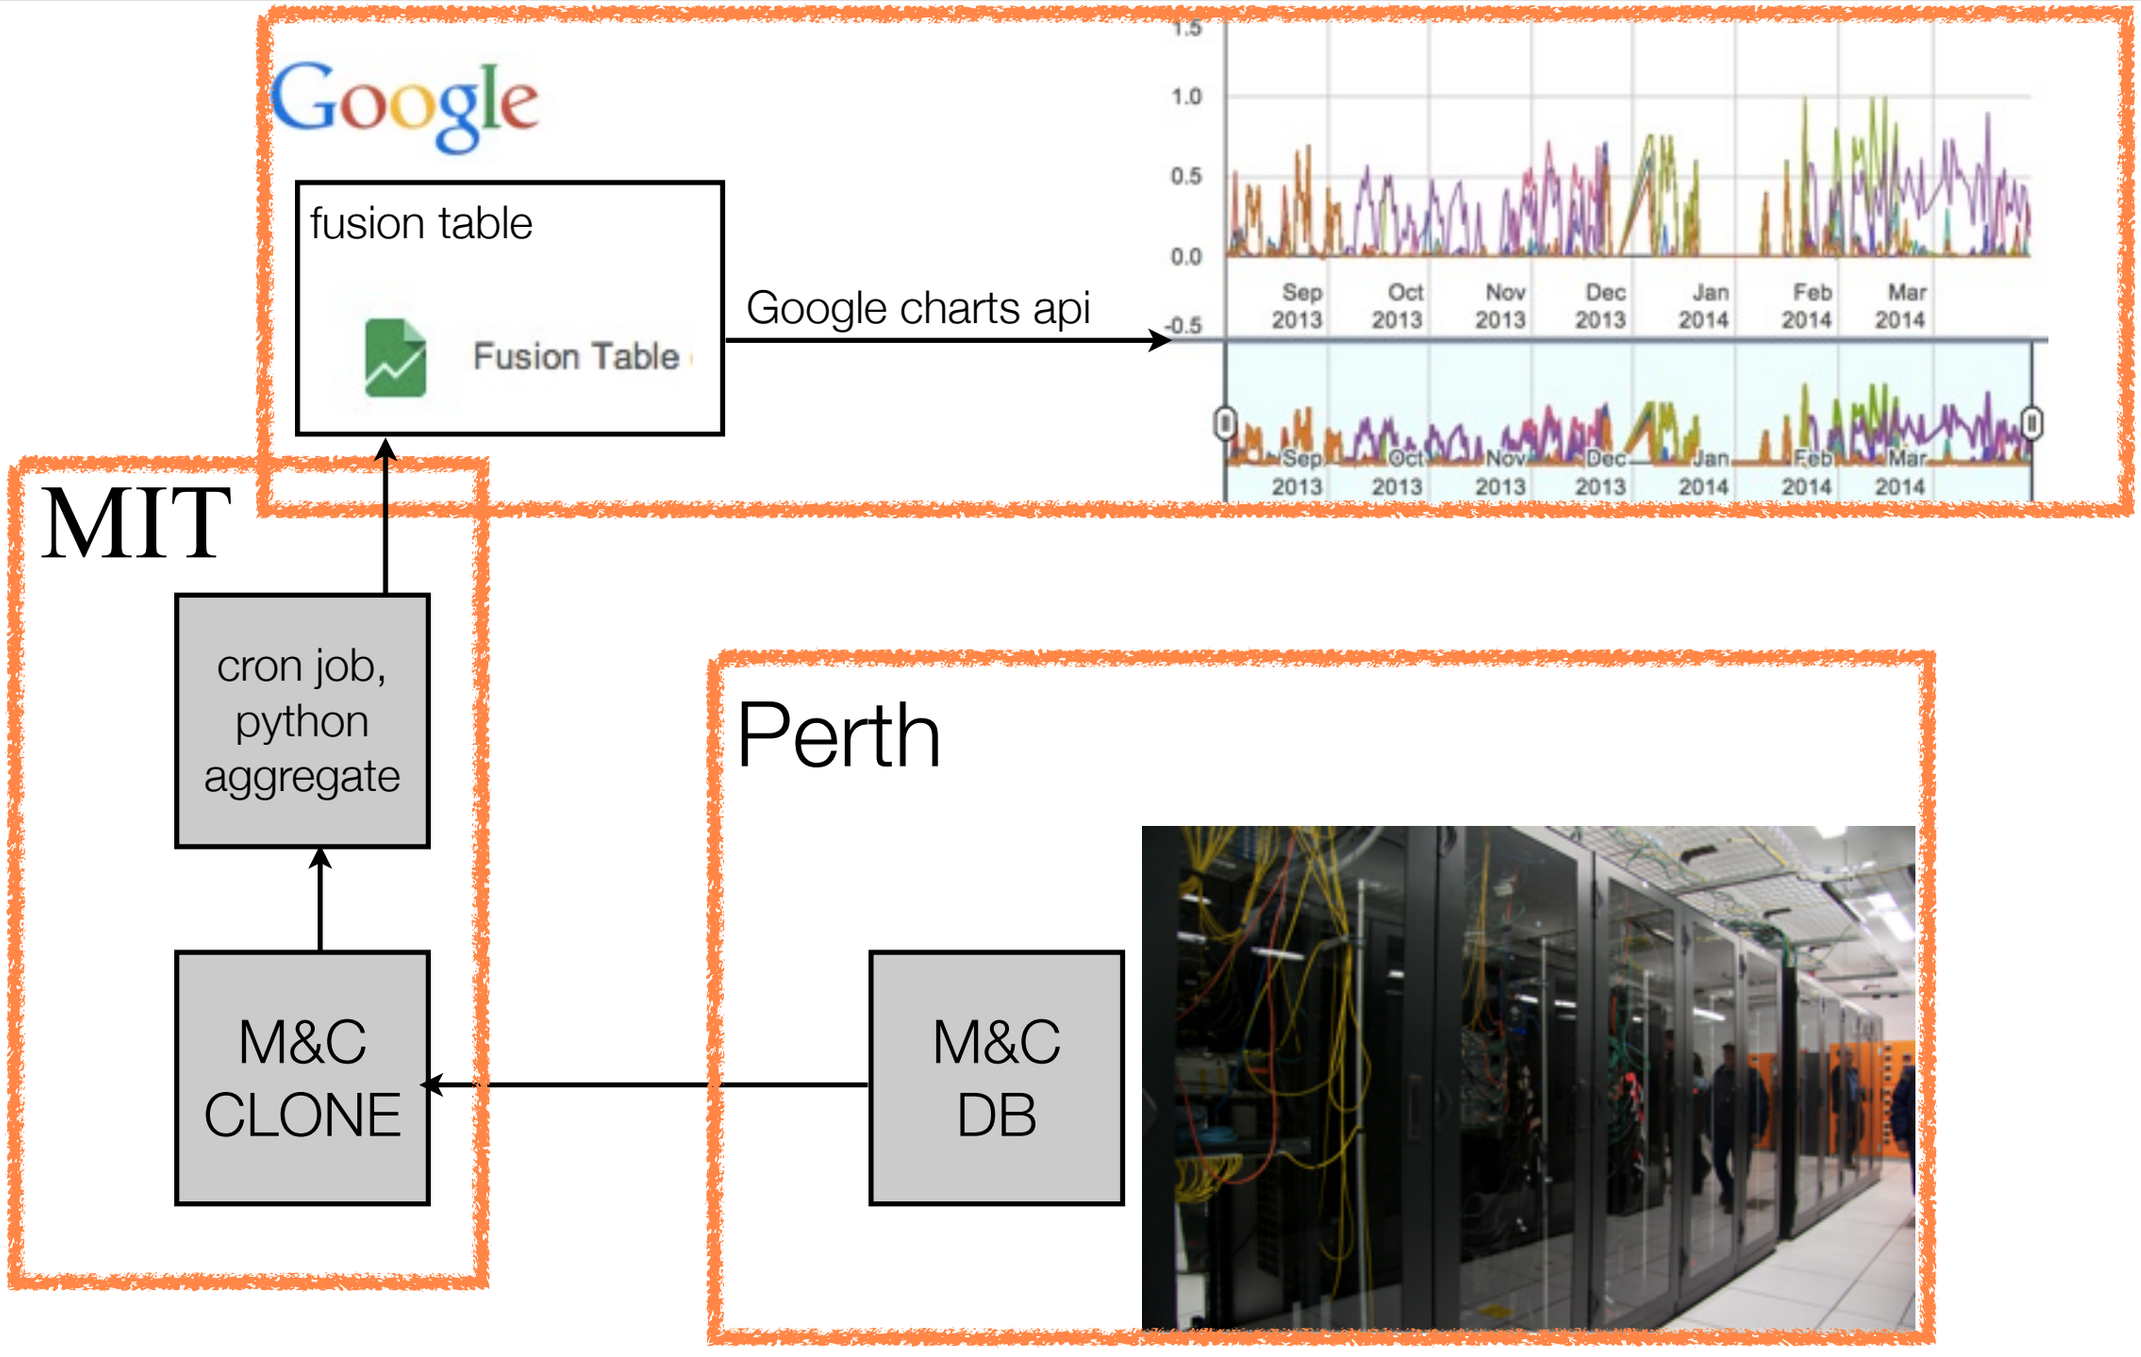

EoR Observing Webpage

Boston	Phoenix	Seattle	Sydney	Perth	Bangalore	UTC
2013-10-26 19:03:16	2013-10-26 16:03:16	2013-10-26 16:03:16	2013-10-27 10:03:16	2013-10-27 07:03:16	2013-10-27 04:33:16	2013-10-26 23:03:16

Useful Links

Observer's Schedule nagios Observation Status Tile Total Power Table of G0009 Observations and their Locations Autocorrelations

Last Update of Day Table Fri Oct 18 20:59:10 UTC 2013 **Last Update of the Observations Table** Fri Oct 18 22:49:26 UTC 2013

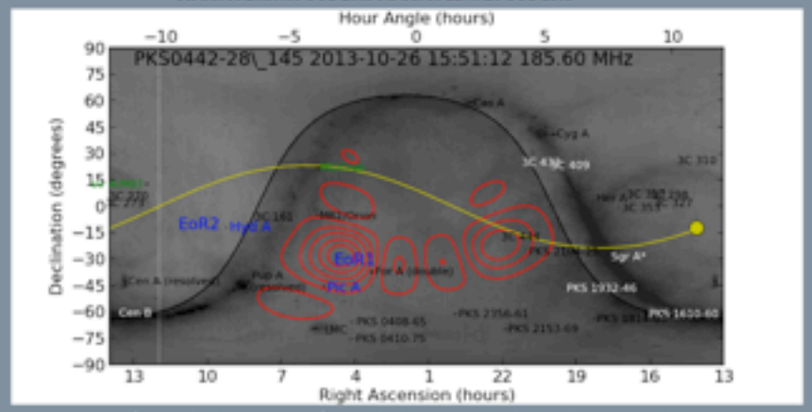

	Observation Number	Observation Name	Project Id	Files	Date
Obs Now	Nothing Now	N/A	N/A	N/A	N/A
Last	1066766168	high_PictorA_season1_2456591	G0009	48	2013-10-25 19:55:51+00
2nd Last	1066766048	high_J2334-4_season1_2456591	G0009	48	2013-10-25 19:53:51+00

Next Observation in : No G0009 Observations Scheduled At this Time.

Next 24 Hours	Total Queued
0	0

Timestamp	Observer Name	Date of observation (UTC)	Notes	Problems
10/25/2013 8:55:08	Ian Sullivan	10/25/2013	Everything looks fine, but tiles 115 and 116 still report they are turned off.	hardware fault
10/23/2013 6:29:02	Aaron Ewall-Wice	10/23/2013	Total Tile Powers and Auto Correlations look fine for now. There is a slight systematic offset in the YY. http://telemetry-static.mwa128t.org/autoplots/1066570138/index.html	Everything is fine
10/21/2013 21:19:59	Andre Offringa	10/9/2013	I imaged 1065388640 (HydA, low), went fine & all unflagged antennae look ok.	Everything is fine
10/21/2013 7:47:07	Danny	10/21/2013	Pawsey center is offline since Friday. Transfers will resume when they resume.	problem with data flow, got an email from ops team
10/17/2013 7:37:07	Danny	10/17/2013	A clock driver has died! There is a spare, but it must be done manually. Observing is cancelled until Monday night.	hardware fault, got an email from ops team
10/15/2013 13:40:04	Ian Sullivan	10/15/2013	End of night summary: Tiles 115 and 116 reported off all night, except for a few minutes. Coarse frequency channel 140 was missing from the autocorellation plots for a while in the middle of the night.	hardware fault
10/15/2013 11:35:34	Ian Sullivan	10/15/2013	No observations scheduled for past half hour, no data recorded. Next observation scheduled in a minute.	No data taken
10/15/2013 9:58:49	Ian Sullivan	10/15/2013	Tiles 115 and 116 are reported powered off again.	hardware fault
10/15/2013 9:58:02	Ian Sullivan	10/15/2013	Coarse frequency channel #140 appears to be missing from the autocorellation plots.	
10/15/2013 9:53:14	Ian Sullivan	10/15/2013	Tiles 115 and 116 now appear fine. Everything else looks good, too!	Everything is fine

Make a Log Entry! View Full Log

Hours of Observing

Receiver Failure Rates

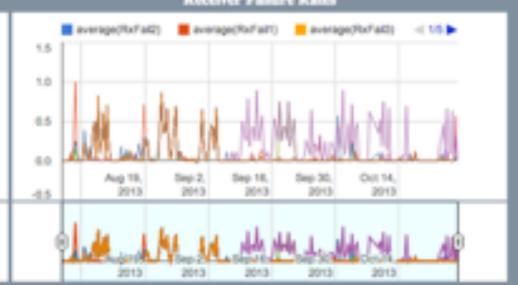

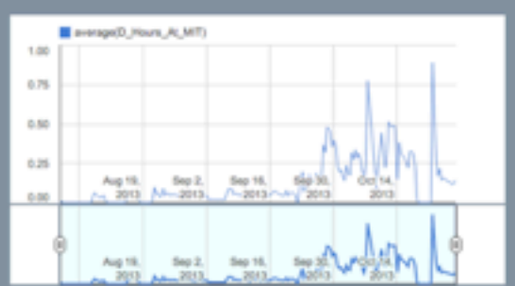

Boston	Phoenix	Seattle	Sydney	Perth	Bangalore	UTC
013-10-26 19:03:16	2013-10-26 16:03:16	2013-10-26 16:03:16	2013-10-27 10:03:16	2013-10-27 07:03:16	2013-10-27 04:33:16	2013-10-26 23:03:16

Useful Links

Observer's Schedule nagios Observation Status Tile Total Power Table of G0009 Observations and their Locations Autocorrelations

Last Update of Day Table **Last Update of the Observations Table**

Fri Oct 18 20:59:10 UTC 2013 Fri Oct 18 22:49:26 UTC 2013

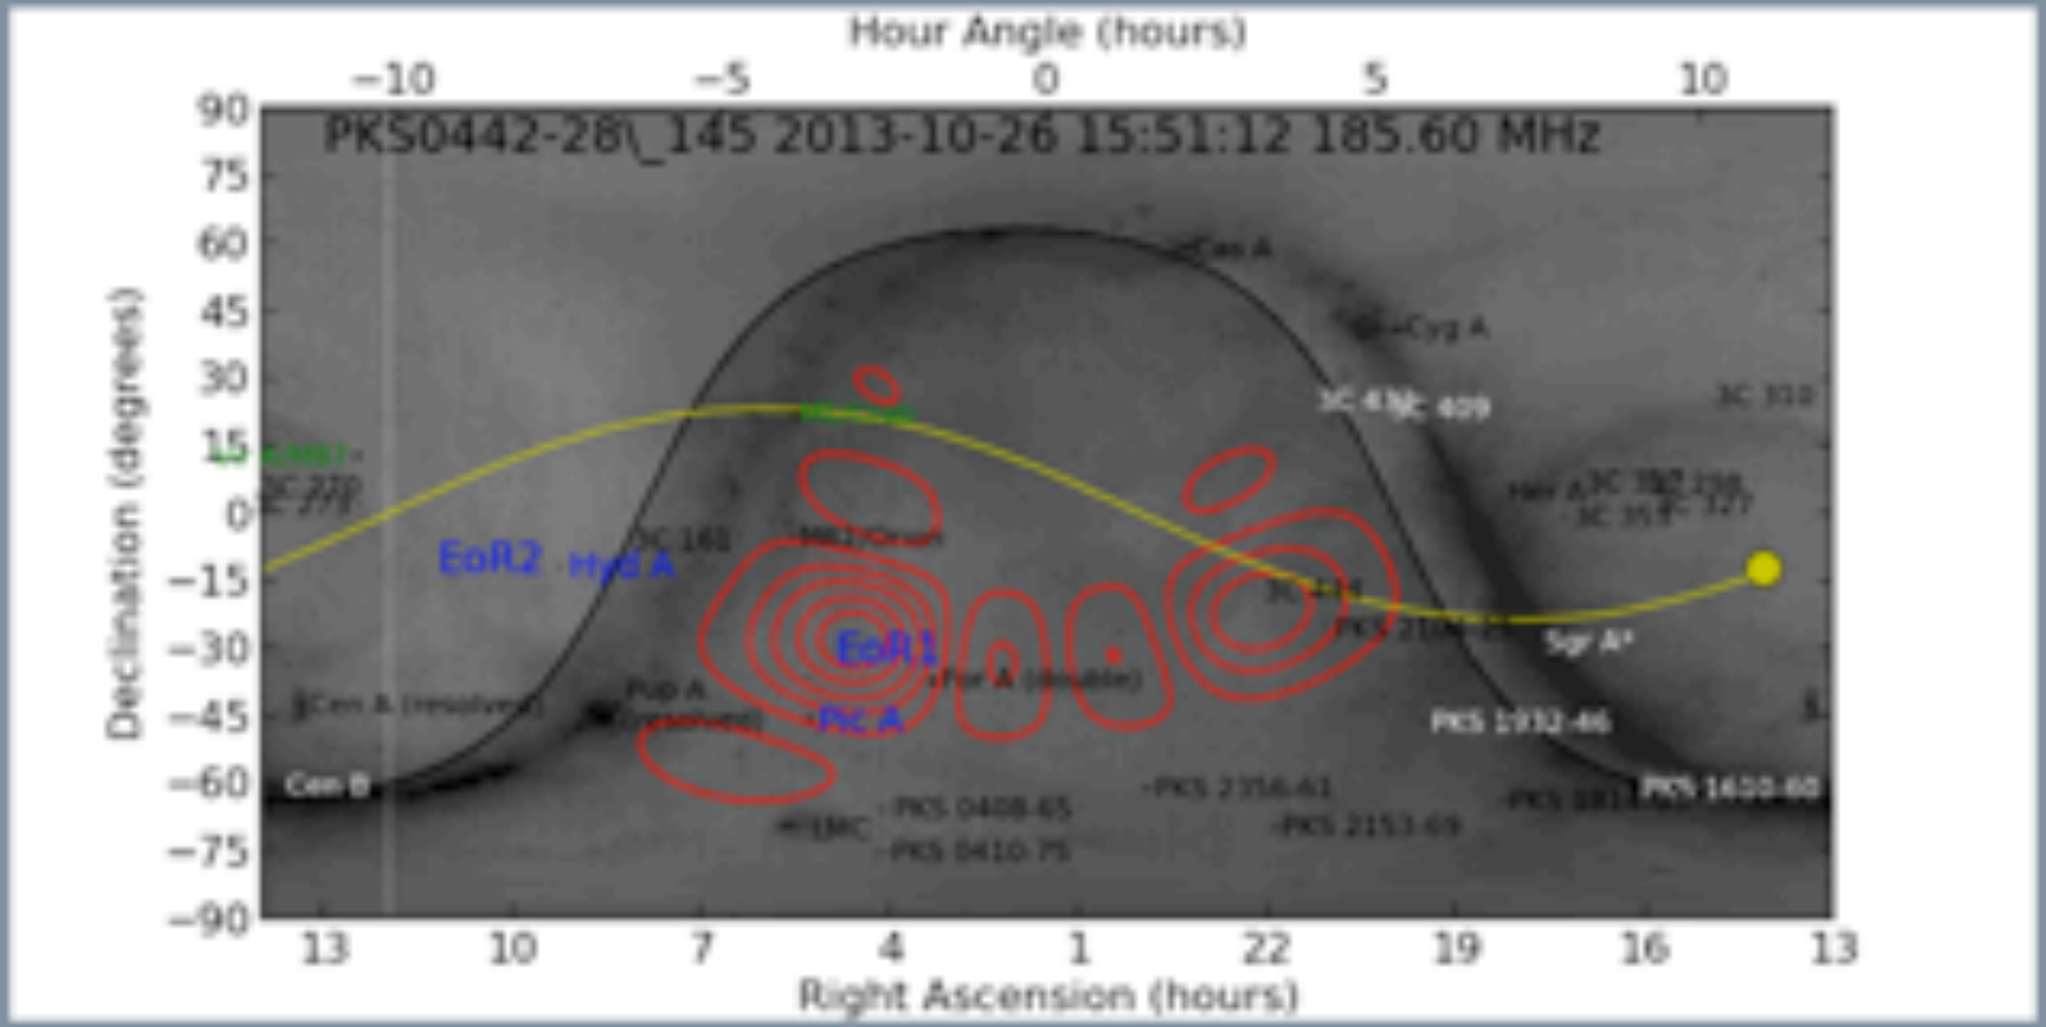

	Observation Number	Observation Name	Project Id	Files	Date
Obs Now	Nothing Now	N/A	N/A	N/A	N/A
Last	1066766168	high_PictorA_season1_2456591	G0009	48	2013-10-25 19:55:51+00
2nd Last	1066766048	high_J2334-4_season1_2456591	G0009	48	2013-10-25 19:53:51+00

Next Observation in : No G0009 Observations Scheduled At this Time.

Next 24 Hours	Total Queued
0	0

Timestamp	Observer Name	Date of observation (UTC)	Notes	Problems

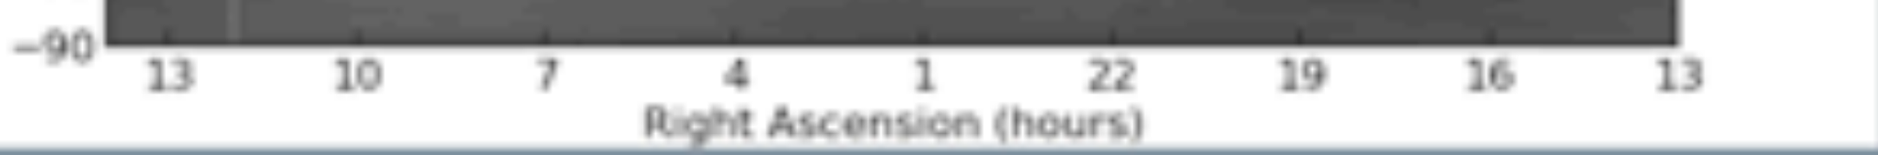

	Observation Number	Observation Name	Project Id	Files	Date
Obs Now	Nothing Now	N/A	N/A	N/A	N/A
Last	1066766168	high_PictorA_season1_2456591	G0009	48	2013-10-25 19:55:51+00
2nd Last	1066766048	high_J2334-4_season1_2456591	G0009	48	2013-10-25 19:53:51+00

Next Observation in : No G0009 Observations Scheduled At this Time.

Next 24 Hours	Total Queued
0	0

Timestamp	Observer Name	Date of observation (UTC)	Notes	Problems
10/25/2013 8:55:08	Ian Sullivan	10/25/2013	Everything looks fine, but tiles 115 and 116 still report they are turned off.	hardware fault
10/23/2013 6:29:02	Aaron Ewall-Wice	10/23/2013	Total Tile Powers and Auto Correlations look fine for now. There is a slight systematic offset in the YY. http://telemetry-static.mwa128t.org/autoplots/1066570138/index.html	Everything is fine
10/21/2013 21:19:59	Andre Offringa	10/9/2013	I imaged 1065388640 (HydA, low), went fine & all unflagged antennae look ok.	Everything is fine
10/21/2013 7:47:07	Danny	10/21/2013	Pawsey center is offline since Friday. Transfers will resume when they resume.	problem with data flow, got an email from ops team
10/17/2013 7:37:07	Danny	10/17/2013	A clock driver has died! There is a spare, but it must be done manually. Observing is cancelled until Monday night.	hardware fault, got an email from ops team
10/15/2013 13:40:04	Ian Sullivan	10/15/2013	End of night summary: Tiles 115 and 116 reported off all night, except for a few minutes. Coarse frequency channel 140 was missing from the autocorellation plots for a while in the middle of the night.	hardware fault
10/15/2013 11:35:34	Ian Sullivan	10/15/2013	No observations scheduled for past half hour, no data recorded. Next observation scheduled in a minute.	No data taken
10/15/2013 9:58:49	Ian Sullivan	10/15/2013	Tiles 115 and 116 are reported powered off again.	hardware fault
10/15/2013 9:58:02	Ian Sullivan	10/15/2013	Coarse frequency channel #140 appears to be missing from the autocorellation plots.	
10/15/2013 9:53:14	Ian Sullivan	10/15/2013	Tiles 115 and 116 now appear fine. Everything else looks good, too!	Everything is fine

fake a Log Entry!

View Full Log

Hours of Observing

Receiver Failure Rates

9:58:49	Sullivan	10/15/2013	Tiles 115 and 116 are reported powered off again.	hardware fault
10/15/2013 9:58:02	Ian Sullivan	10/15/2013	Coarse frequency channel #140 appears to be missing from the autocorellation plots.	
10/15/2013 9:53:14	Ian Sullivan	10/15/2013	Tiles 115 and 116 now appear fine. Everything else looks good, too!	Everything is fine

fake a Log Entry!

View Full Log

Hours of Observing

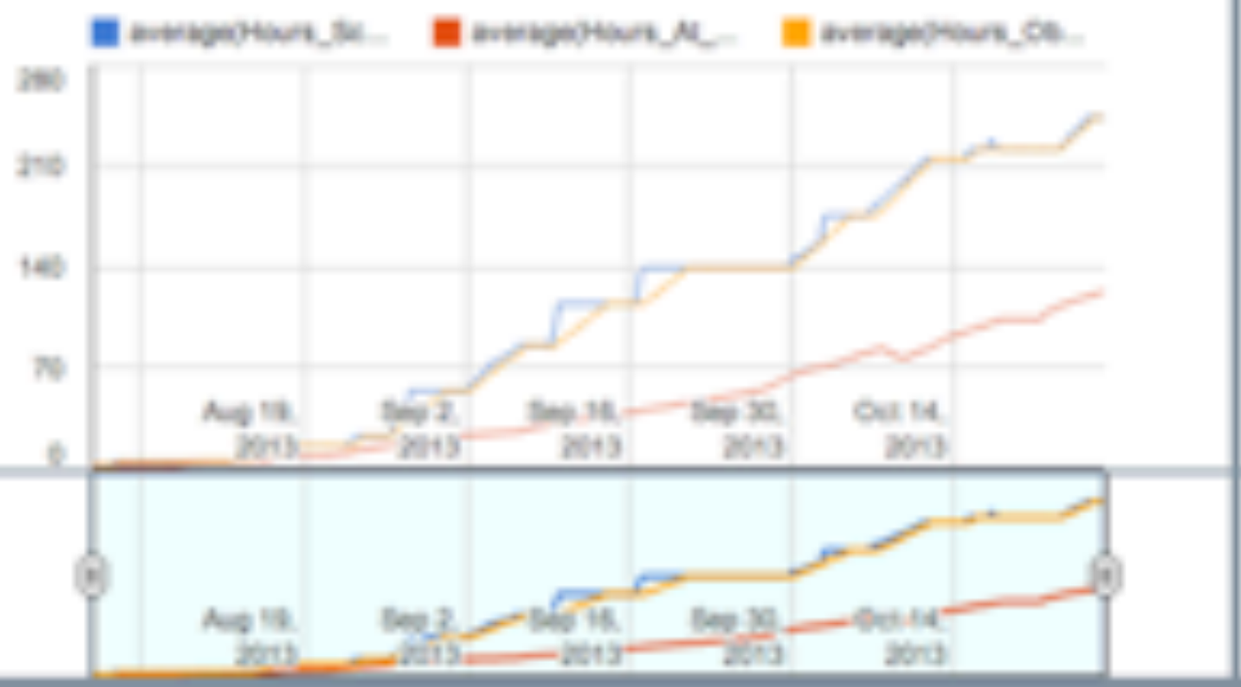

Receiver Failure Rates

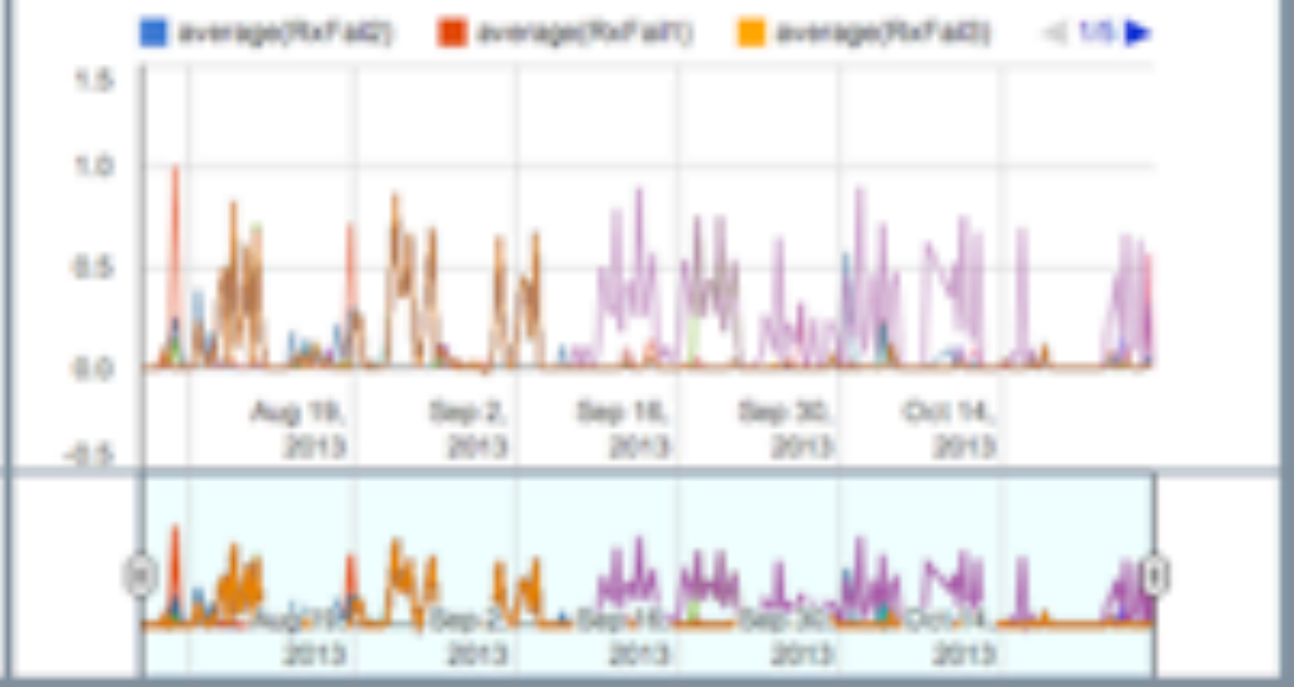

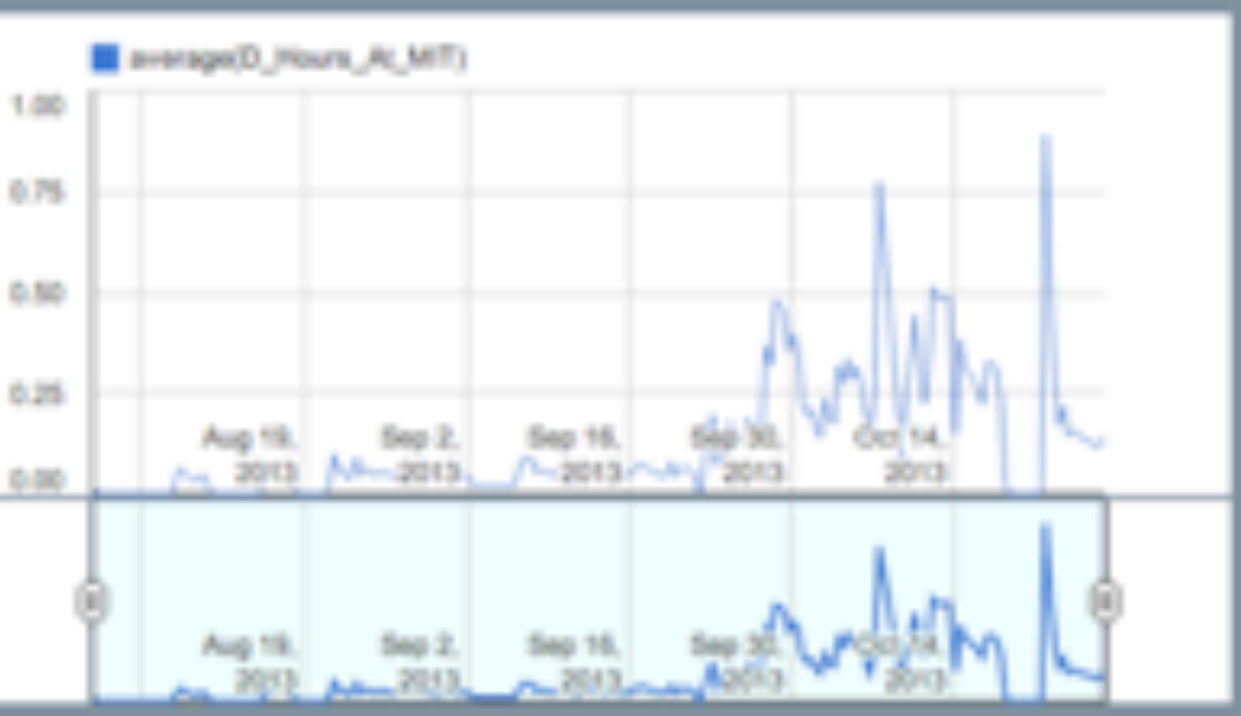

9:58:49	Sullivan	10/15/2013	Tiles 115 and 116 are reported powered off again.
10/15/2013 9:58:02	Ian Sullivan	10/15/2013	Coarse frequency channel #140 appears to be missing from the autocorellation plots.
10/15/2013 9:53:14	Ian Sullivan	10/15/2013	Tiles 115 and 116 now appear fine. Everything else looks good, too!

Make a Log Entry!

View Full Log

Hours of Observing

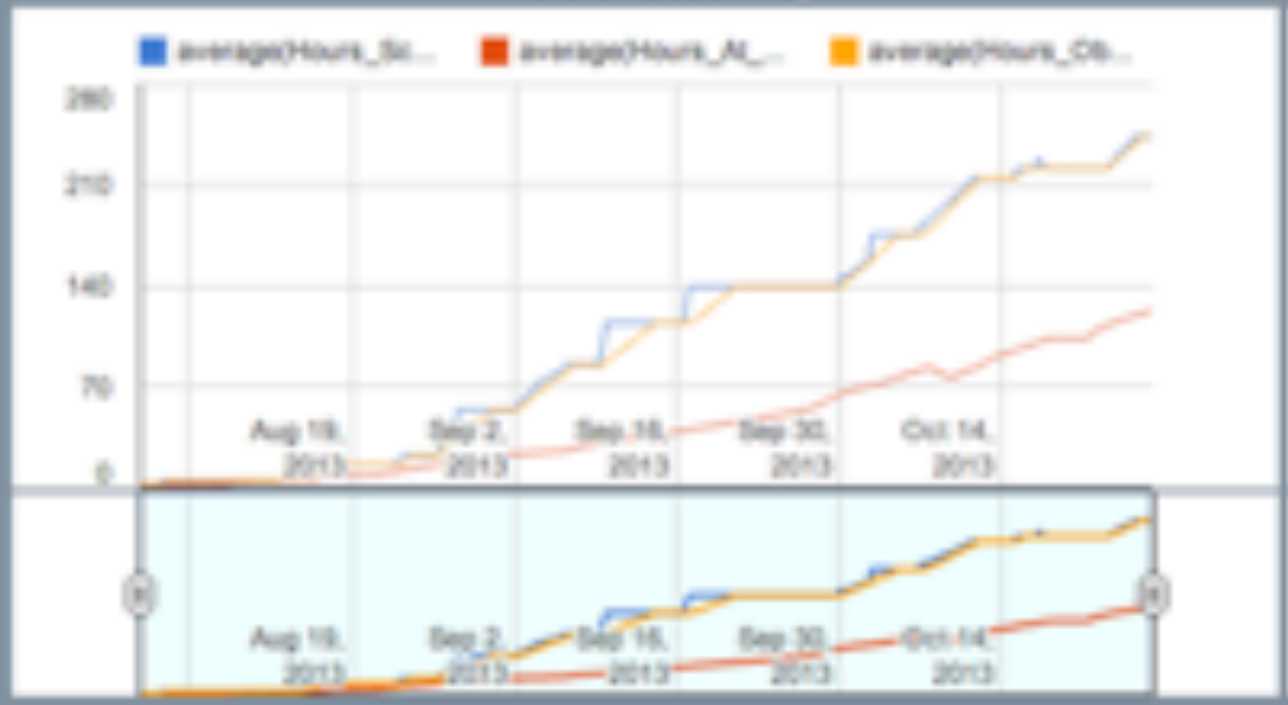

Receiver Failure Rates

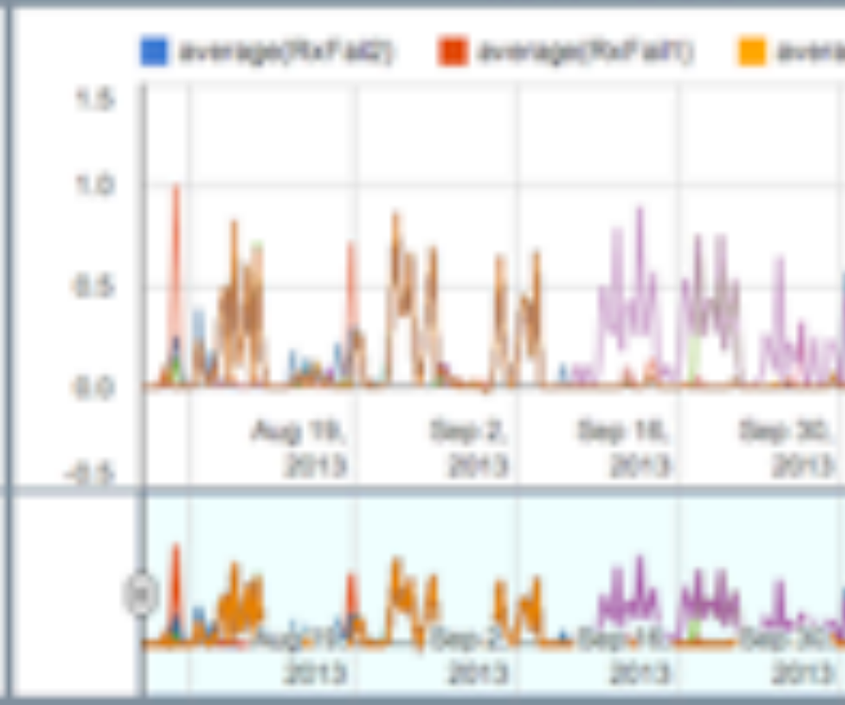

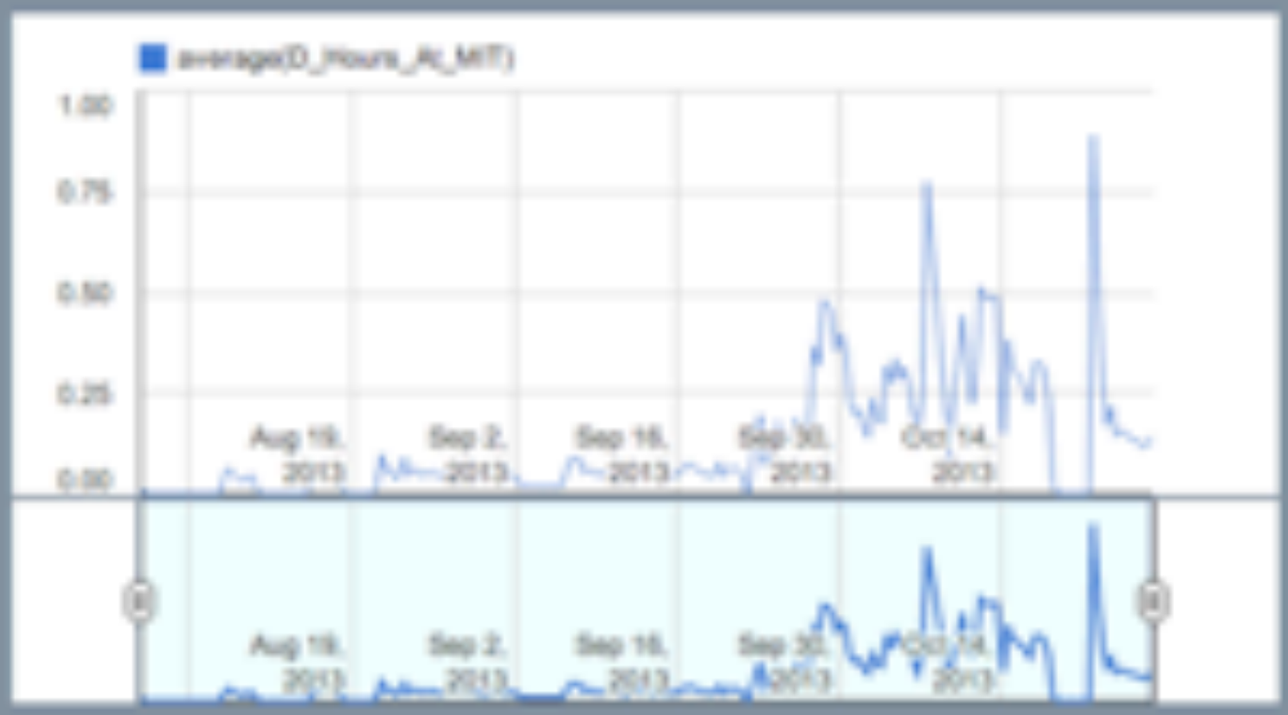

PAPER stats

Saturday, 22 March 2014

Daily Data Report

--- Transfer summary from paper1 ---

--- Free space on various machines ---
pot0:
/dev/sdb1 55T 53T 2.3T 96% /data0
!!!BAD!!!
paper1:
/dev/sda1 3.7T 3.1T 614G 84% /data0
/dev/sdb1 3.7T 3.7T 11G 100% /data1

--- Lost Lambs? ---
Latest mod to pot0 is on /data0/psa6697
0 missing files.
Latest mod to pot1 is on
0 missing files.

This daily report is located in qmaster:/home/obs/AutomaticallyGeneratedReports/2014-03-22.log

Posted by Team PAPER at 11:00 No comments:

+1 Recommend this on Google

Friday, 21 March 2014

Daily Data Report

--- Transfer summary from paper1 ---

--- Free space on various machines ---
pot0:
/dev/sdb1 55T 53T 2.3T 96% /data0
!!!BAD!!!
paper1:
/dev/sda1 3.7T 3.1T 614G 84% /data0
/dev/sdb1 3.7T 3.7T 11G 100% /data1

--- Lost Lambs? ---
Latest mod to pot0 is on /data0/psa6697
0 missing files.
Latest mod to pot1 is on
0 missing files.

This daily report is located in qmaster:/home/obs/AutomaticallyGeneratedReports/2014-03-21.log

Followers

Join this site
with Google Friend Connect

Members (1)

Already a member? Sign in

Blog Archive

rm /data1/obs/2456723/trigger_2456723.txt.sent
rmdir /data1/obs/2456723

Free space:

Filesystem Size Used Avail Use% Mounted on
/dev/disk/by-uuid/79e30170-a20a-4a03-919e-cff8260772ff 112G 53G 54G 50% /
/dev/sda1 3.7T 2.1T 1.7T 56% /data0
/dev/sdb1 3.7T 2.1T 1.7T 56% /data1

Posted by Team PAPER at 08:00 No comments:

Recommend this on Google

Monday, 10 March 2014

[ALERT] Insufficent space on pot0:/data0

Data for JD 2456727 requires 1086580516+10000 KB, but
pot0:/data0 has only 191581480 KB available.

Data transfer for JD 2456727 aborting!

Posted by Team PAPER at 21:10 No comments:

Recommend this on Google

Newer Posts Home Older Posts

Subscribe to: Posts (Atom)

Simple template. Powered by Blogger.

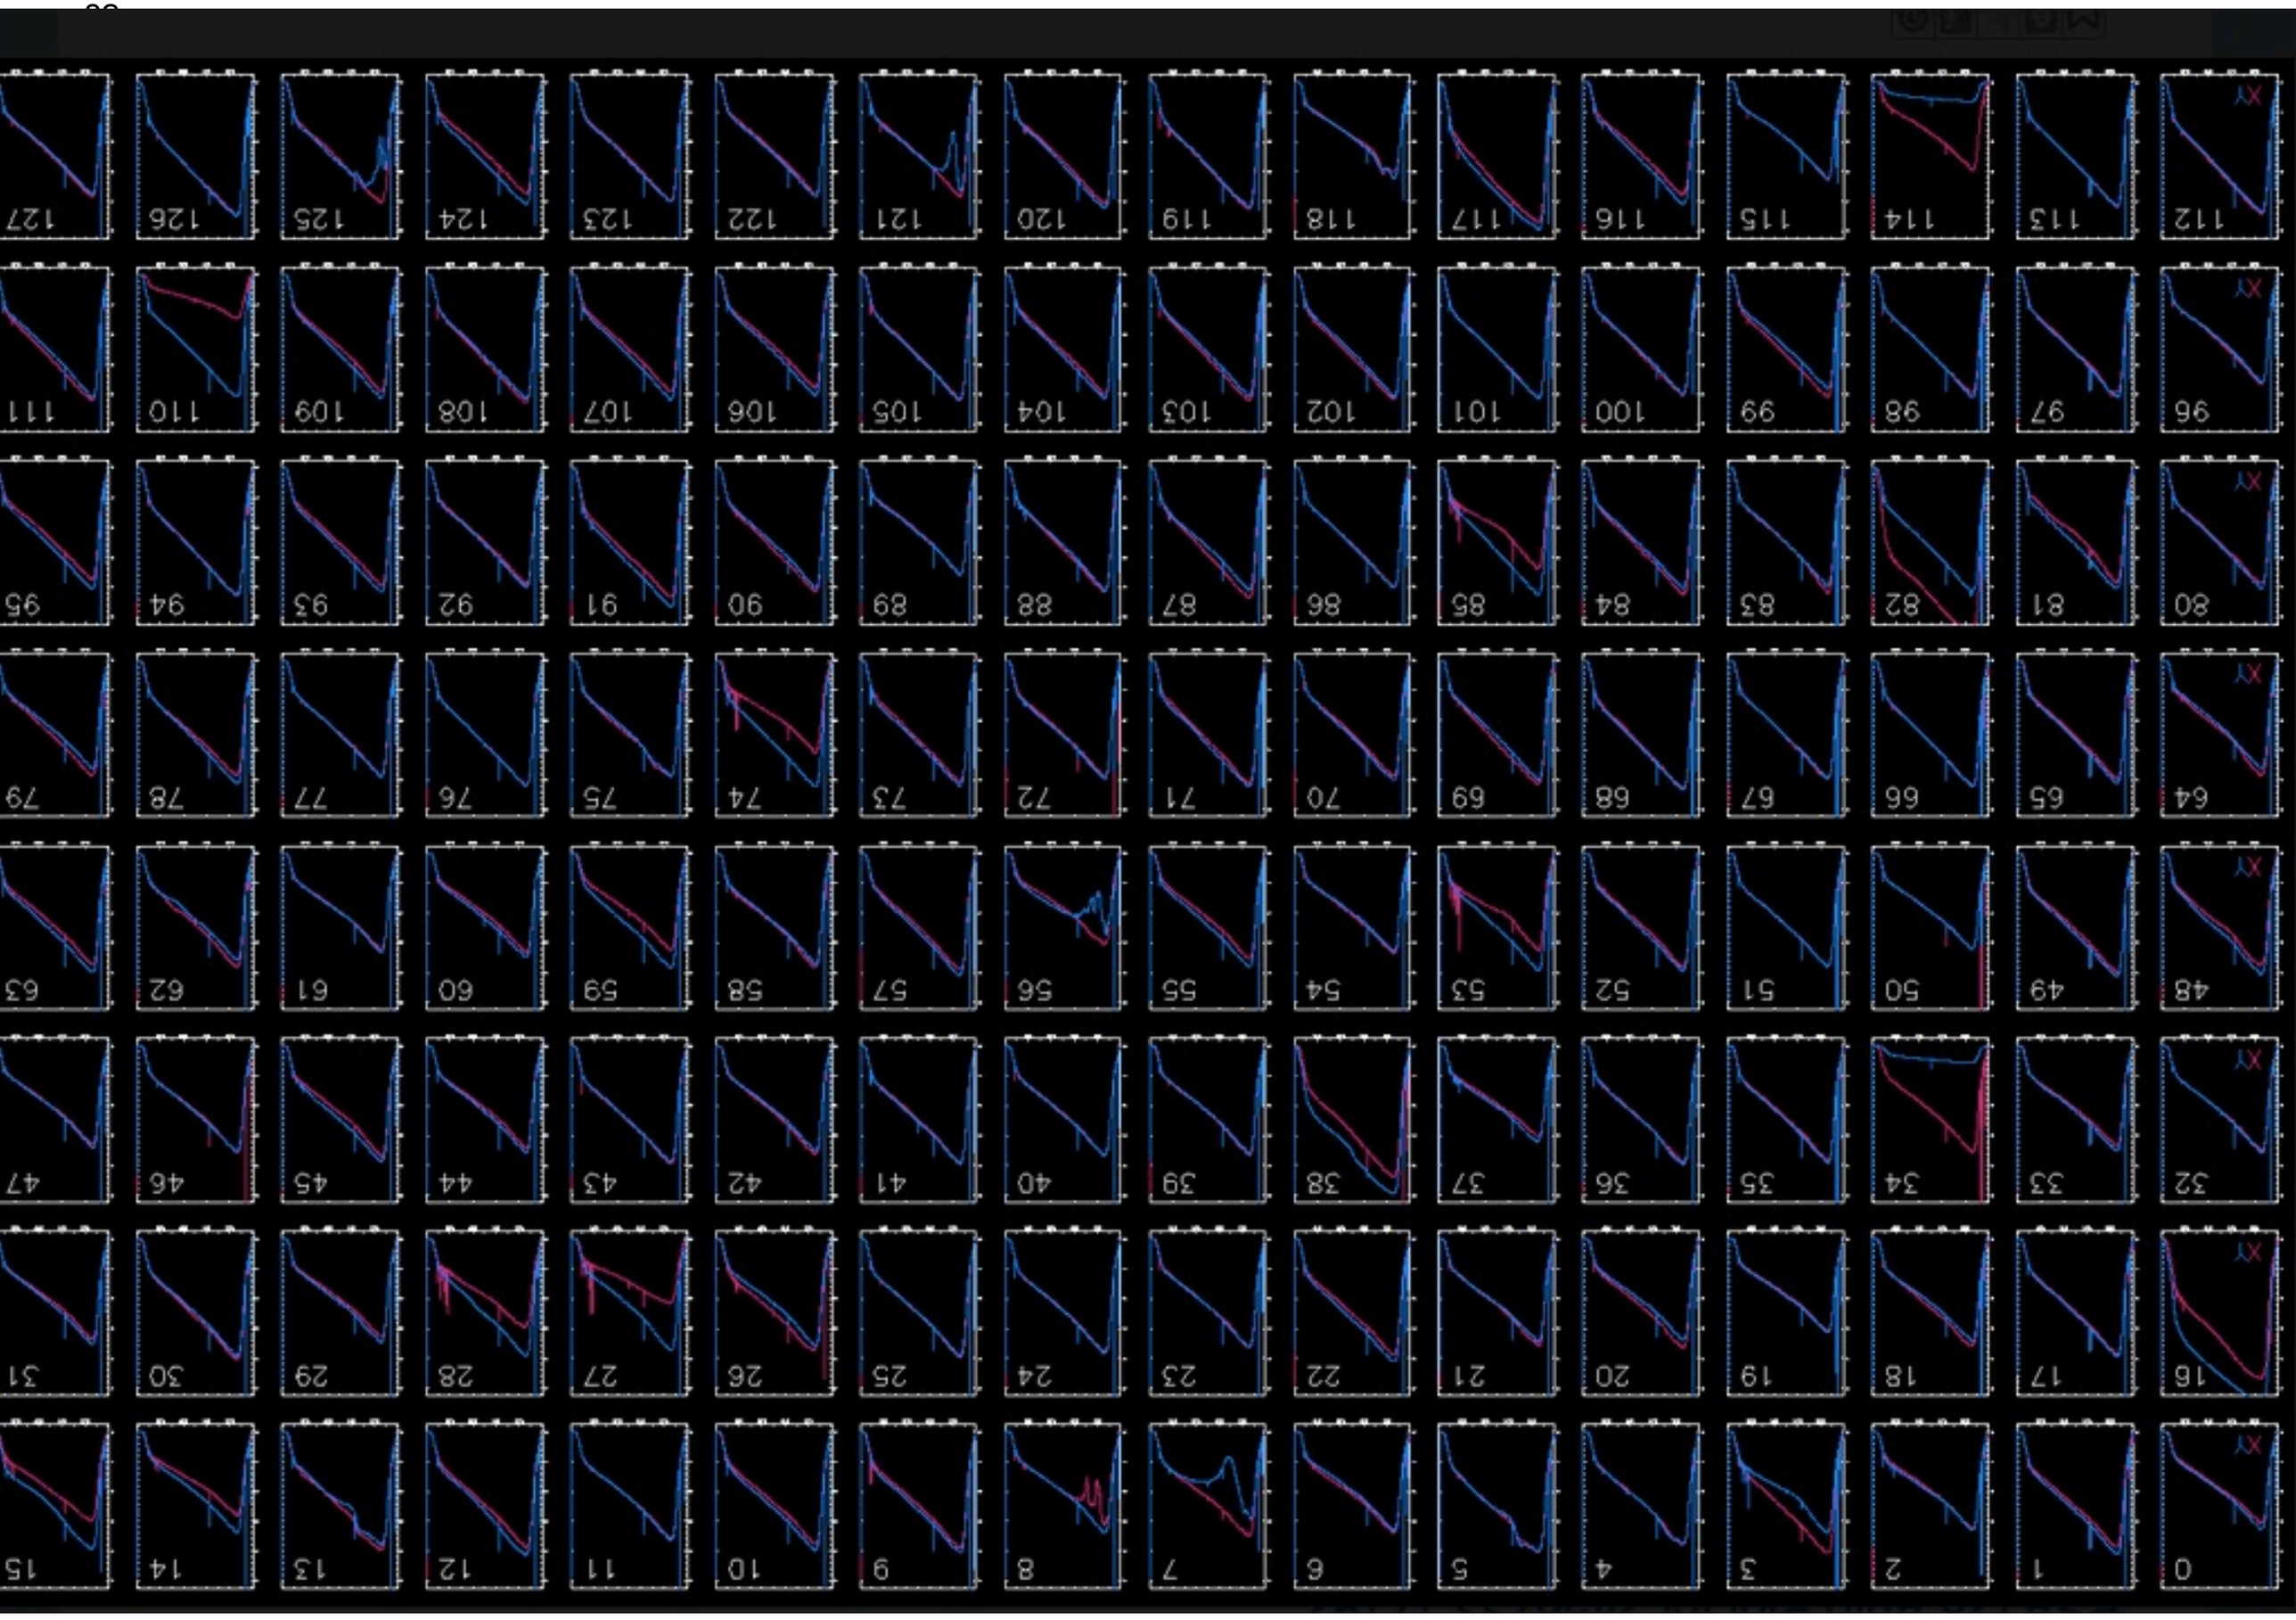
127 126 125 124 123 122 121 120 119 118 117 116 115 114 113 112
111 110 109 108 107 106 105 104 103 102 101 100 99 98 97 96
95 94 93 92 91 90 89 88 87 86 85 84 83 82 81 80
79 78 77 76 75 74 73 72 71 70 69 68 67 66 65 64
63 62 61 60 59 58 57 56 55 54 53 52 51 50 49 48
47 46 45 44 43 42 41 40 39 38 37 36 35 34 33 32
31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16
15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0

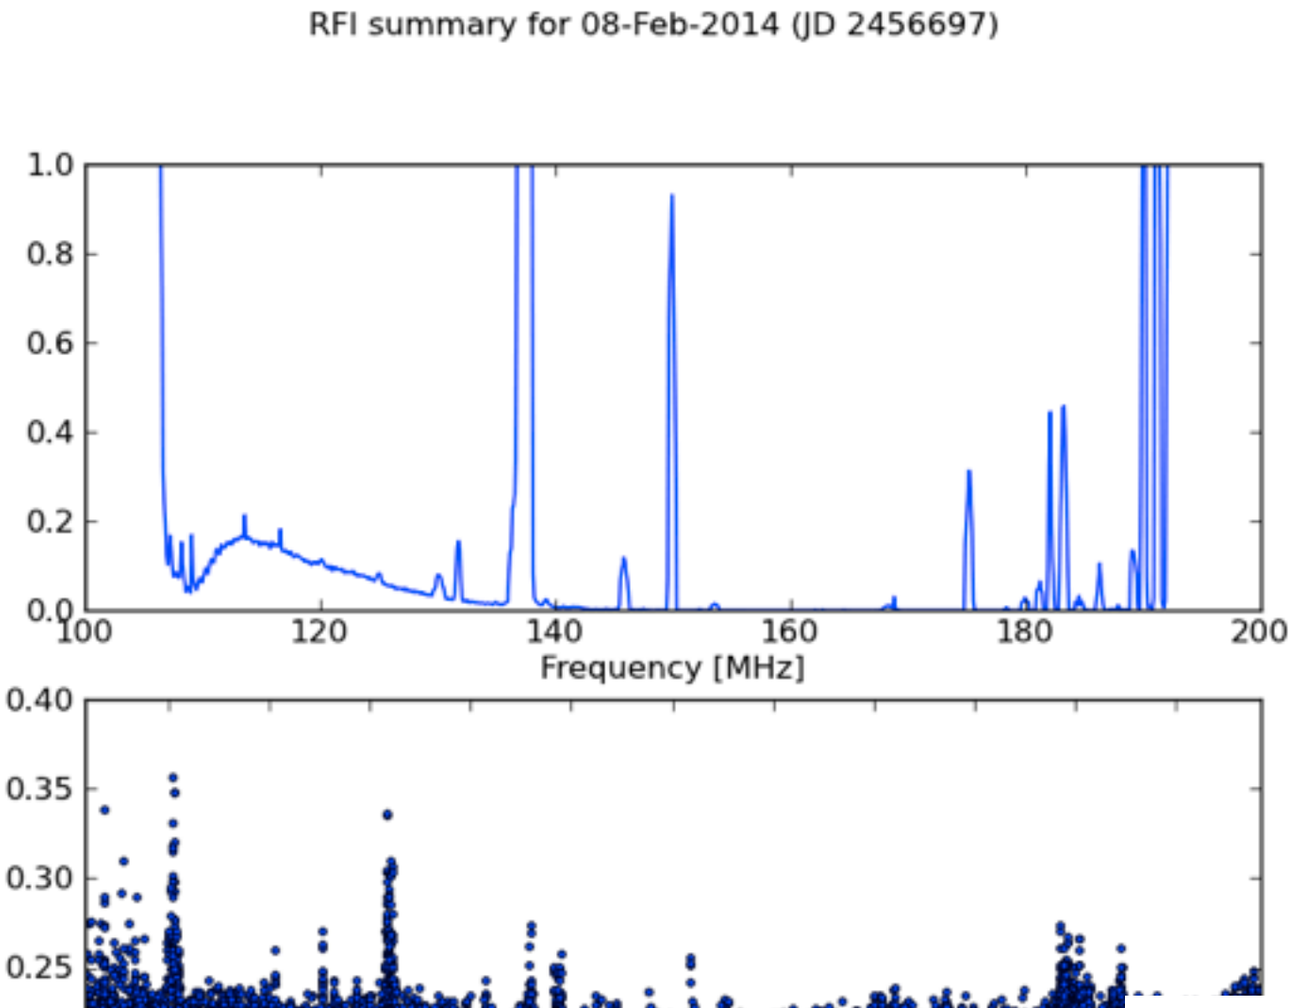
RFI summary for 08-Feb-2014 (JD 2456697)
Frequency [MHz]
SAST [hours]

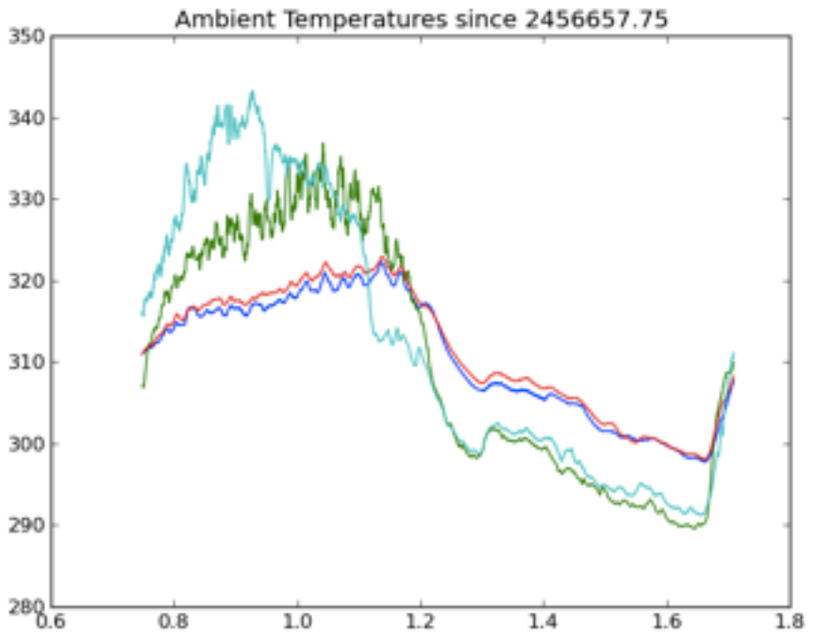
Ambient Temperatures since 2456657.75

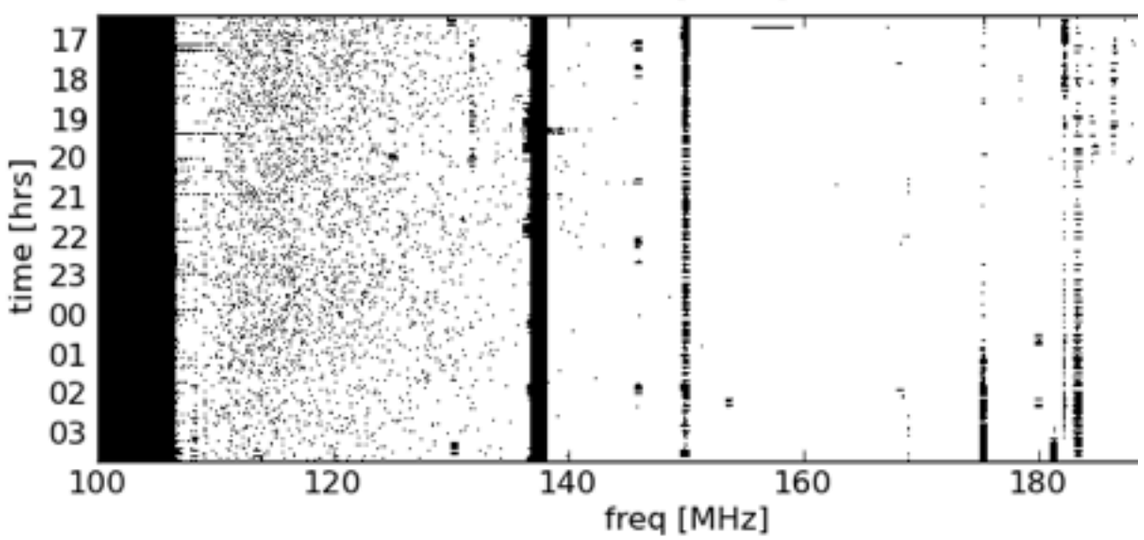
time [hrs]
freq [MHz]

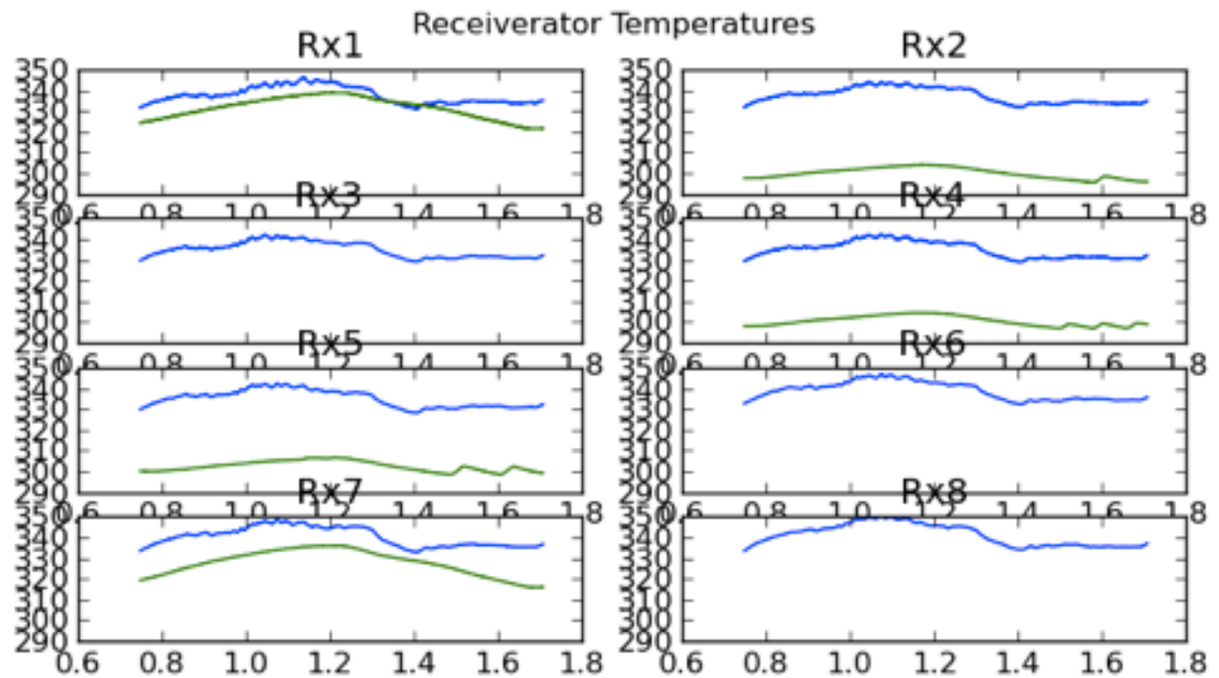
Receiverator Temperatures
Rx1
Rx2
Rx3
Rx4
Rx5
Rx6
Rx7
Rx8

EoRLive status dashboard

Hydrogen Epoch of Reionization Array (HERA)

D. Jacobs **Sign out - Settings**

Array - Network - Data Analysis - Forum

New Forum post
Z. Smith (student):
Check out the image stream from yesterday! Why does that one look different?

Array Status Browser

Array Status Summary: 98% up, 2%

expand +

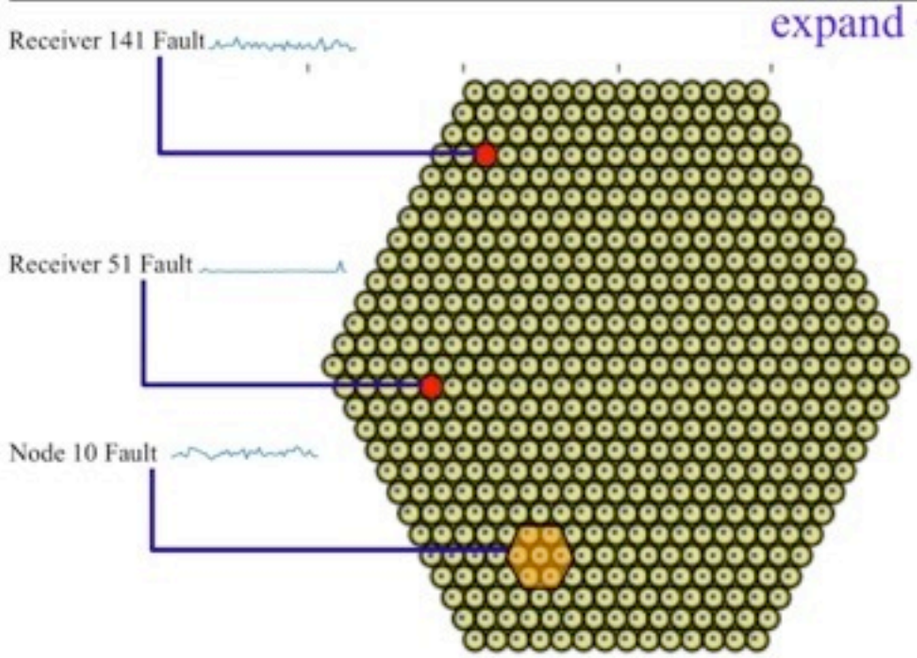

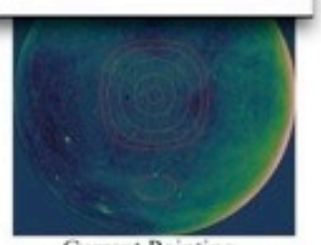

Current Pointing
LST=01h13m

Recent Logs
08/10/2015 - Node 11 Fault flagged - D. DeBoer
08/09/2015 - RFI warning from site, check data! - M. Morales

Observers Online

- D. Jacobs
- A. Marble (student)
- A. Parsons
- M. Morales
- Z. Smith (student)

Flag Categories

fail: 12%
interference: 9%
missing data: 15%
poor calibration: 2%

Network Status

- HERA
- Karoo-CPT
- CPT - UPenn
- UPenn - MIT

Data Processing Status Browser

Quick look Imaging 08/09/2015 expand +

Quick look power spectra 08/09/2015

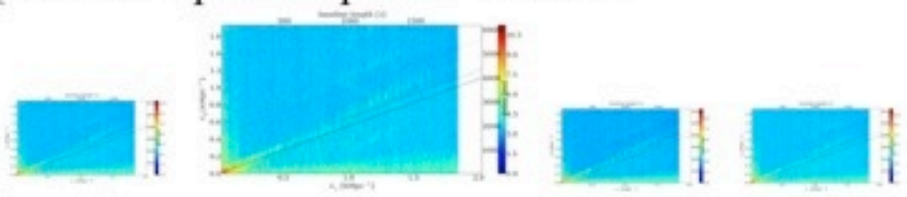

Recent Logs
08/09/2015 - Re: RFI warning from site, check data!
...Power spectra look fine A. Parsons
08/06/2015 - sum fuzz in the high delay bins - D. Jacobs

EoRLive status dashboard

Hydrogen Epoch of Reionization Array (HERA)

D. Jacobs **Sign out** - **Settings**

Array - Network - Data Analysis - Forum

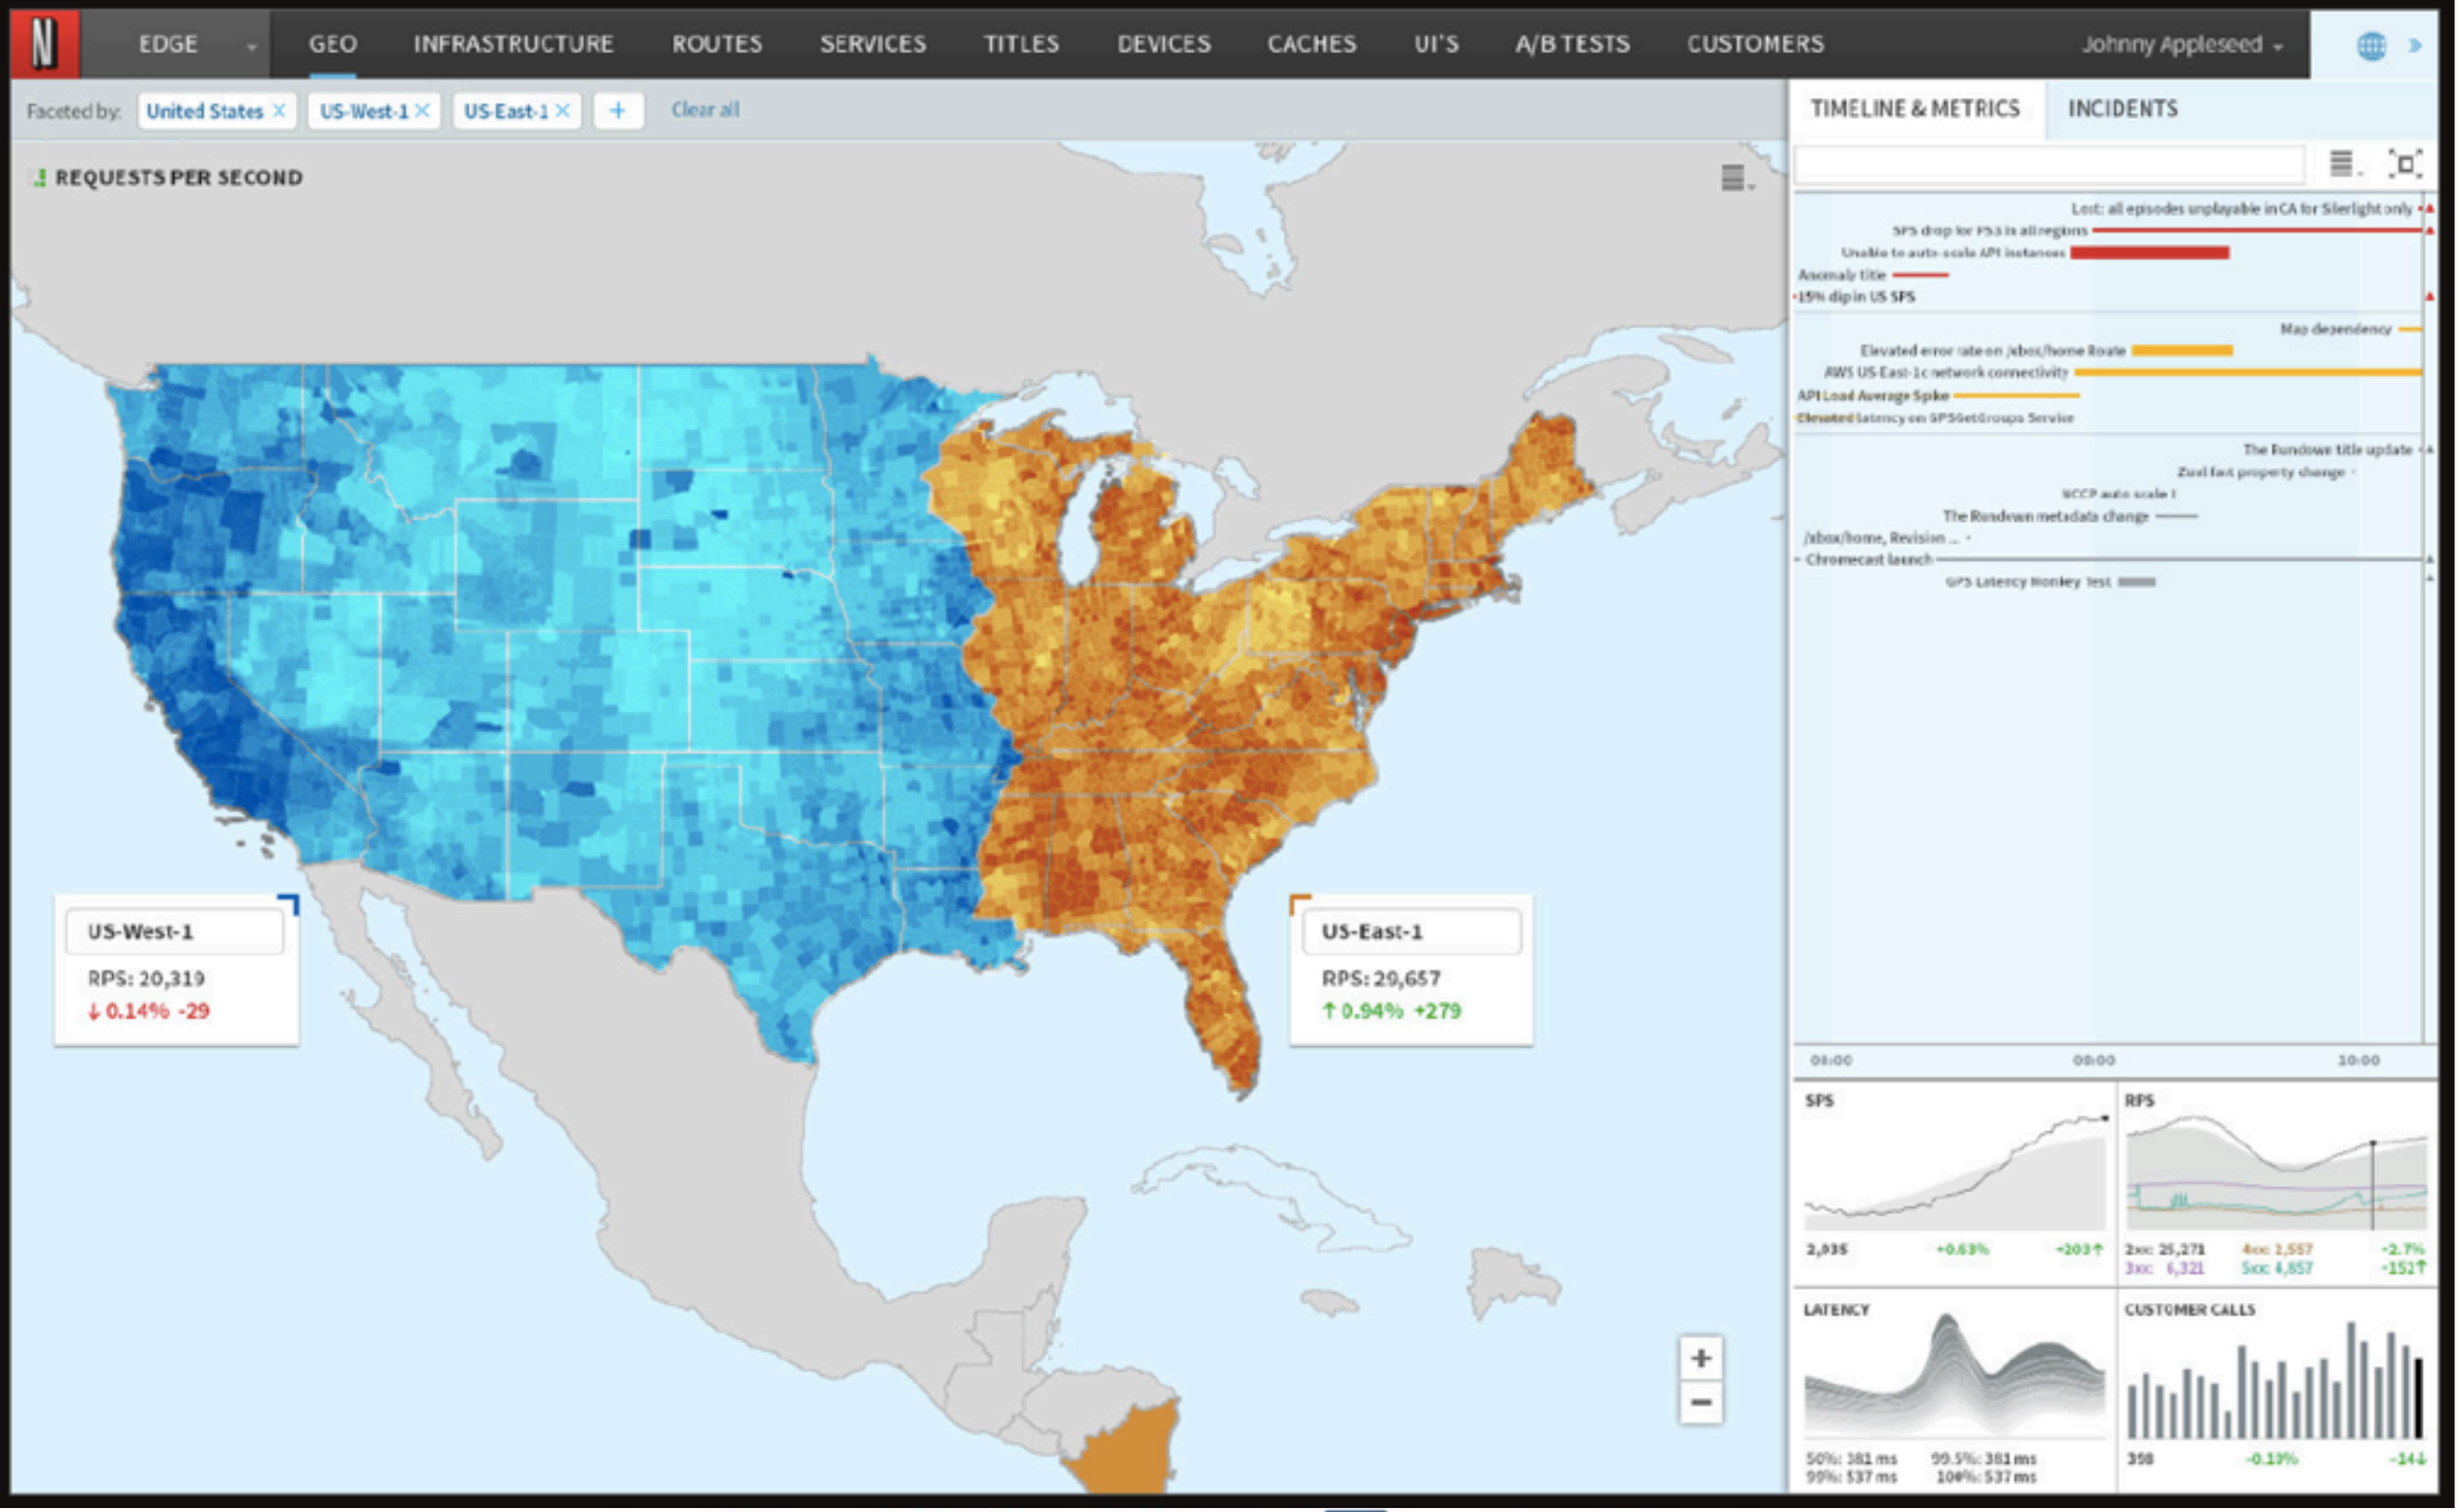

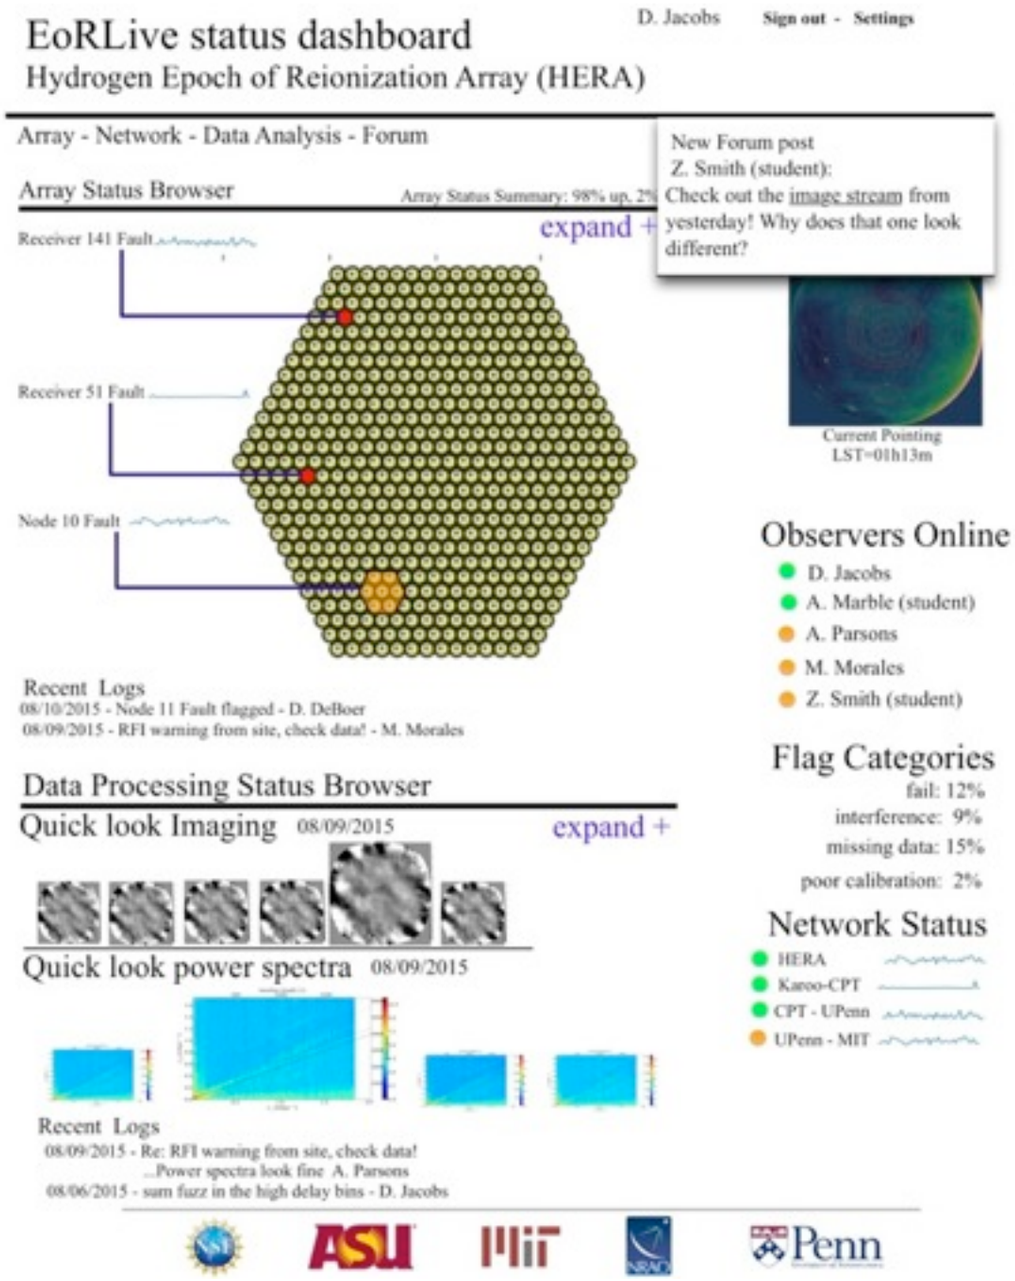
EoRLive status dashboard
Hydrogen Epoch of Reionization Array (HERA)
D. Jacobs
Sign out - Settings
Array - Network - Data Analysis - Forum
New Forum post
Z. Smith (student):
Check out the image stream from yesterday! Why does that one look different?
Array Status Browser
Array Status Summary: 98% up, 2%
expand +
Receiver 141 Fault
Receiver 51 Fault
Node 10 Fault
Current Pointing
LST=01h13m
Observers Online
D. Jacobs
A. Marble (student)
A. Parsons
M. Morales
Z. Smith (student)
Recent Logs
08/10/2015 - Node 11 Fault flagged - D. DeBoer
08/09/2015 - RFI warning from site, check data! - M. Morales
Flag Categories
fail: 12%
interference: 9%
missing data: 15%
poor calibration: 2%
Data Processing Status Browser
Quick look Imaging 08/09/2015
expand +
Network Status
HERA
Karoo-CPT
CPT - UPenn
UPenn - MIT
Quick look power spectra 08/09/2015
Recent Logs
08/09/2015 - Re: RFI warning from site, check data!
...Power spectra look fine A. Parsons
08/06/2015 - sum fuzz in the high delay bins - D. Jacobs
ASU
MIT
Penn

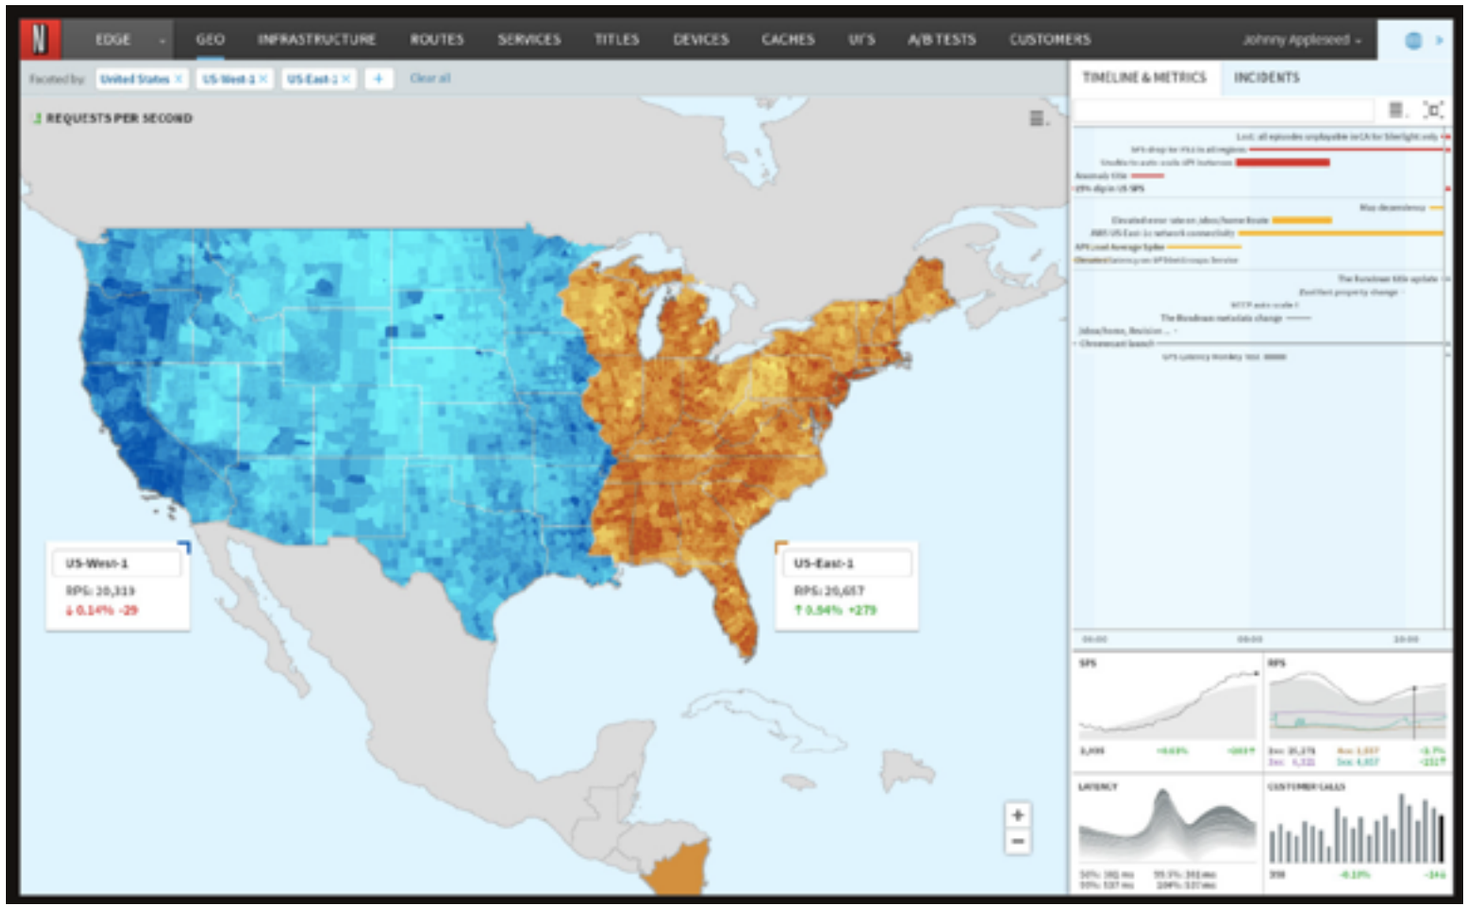
EDGE
GEO
INFRASTRUCTURE
ROUTES
SERVICES
TITLES
DEVICES
CACHES
UI'S
A/B TESTS
CUSTOMERS
Johnny Appleseed
TIMELINE & METRICS
INCIDENTS
REQUESTS PER SECOND
US-West-1
US-East-1

Transitioning to HPC

Transitioning to HPC

- apparent storage cost

Transitioning to HPC

- apparent storage cost
- user custom

Transitioning to HPC

- apparent storage cost
- user custom
- we don’t know what we are doing

Transitioning to HPC

- apparent storage cost
- user custom
- we don’t know what we are doing
  - -> the blind bucket brigade problem has emerged twice in the past year

Transitioning to HPC

- apparent storage cost
- user custom
- we don’t know what we are doing
  - -> the blind bucket brigade problem has emerged twice in the past year


350 GBps
20,662 km
North Pacific Ocean
United States
Mexico
Gulf of California
Gulf of Mexico
Caribbean Sea
Cuba
Puerto Rico
Guatemala
Nicaragua
Venezuela
Guyana
Colombia
Ecuador
Peru
Brazil
Sea of Japan
South Korea
Japan
East China Sea
South China Sea
Vietnam
Philippines
Indonesia
Banda Sea
Arafura Sea
Papua New Guinea
Solomon Sea
Coral Sea
Australia
Great Australian Bight
Tasman Sea
WA
NT
QLD
SA
NSW
VIC
ACT
TAS
ON
QC
NB
PE
NS
ME
VT
NH
MA
RI
CT
NY
PA
NJ
DE
MD
VA
WV
KY
OH
IN
IL
MI
WI
MN
SD
ND
MT
OR
ID
WY
NV
UT
CO
AZ
NM
CA
NE
KS
OK
TX
MO
AR
LA
MS
AL
GA
FL
SC
NC
TN
HI

hours per hour
~TB/hour
1.8
1.2
0.6
0.0
-0.6
Sep 2013
Oct 2013
Nov 2013
Dec 2013
Jan 2014
Feb 2014
Mar 2014

Data flow

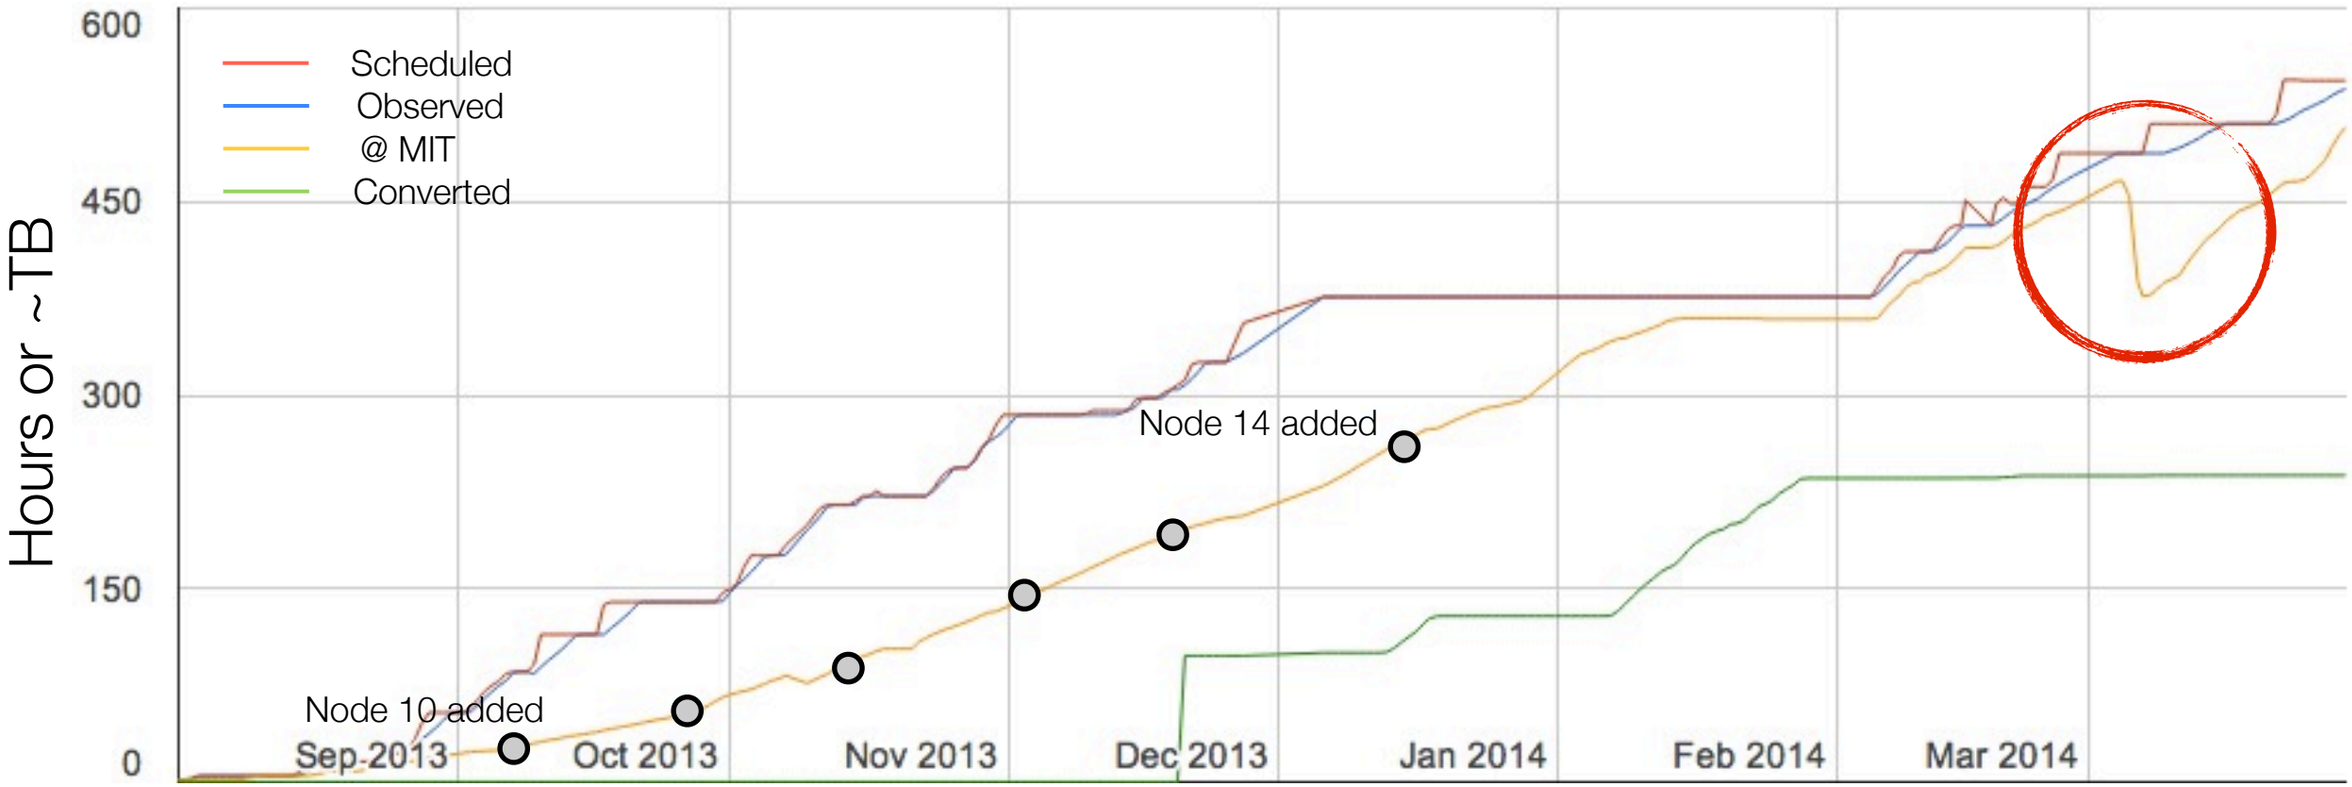
Hours or ~TB
600
450
300
150
0
Scheduled
Observed
@ MIT
Converted
Node 10 added
Node 14 added
Sep 2013
Oct 2013
Nov 2013
Dec 2013
Jan 2014
Feb 2014
Mar 2014

MWA Processing times

Step	Times
record observation	1.86 minute
transfer to Primary Archive in Perth	20 seconds
transfer to US MIT EoR Archive	45 seconds
Average, flag, convert	12 minutes
calibrate, subtract foreground mode, grid	45 minutes
sidereal time average	~1 minute

MWA times

Summary:
total linear time to process 1000 hour observation = 3.66 years
estimated wall time ~ 47 days

Step	Times
record observation	1.86 minute
transfer to Primary Archive in Perth	20 seconds
transfer to US MIT EoR Archive	45 seconds
Average, flag, convert	12 minutes
calibrate, subtract foreground mode, grid	45 minutes
sidereal time average	~1 minute

Computing on the data

- Shared resource with different usage models
  - transferring data
    (newly arriving data, hashing, results of averaging)
  - converting data
    (data bound, uses many cores)
  - subtract and average
    (data bound)
  - lst average
    (combines many data points, a corner turn)
  - other cpu intensive jobs (data archive is also largest and most capable machine)

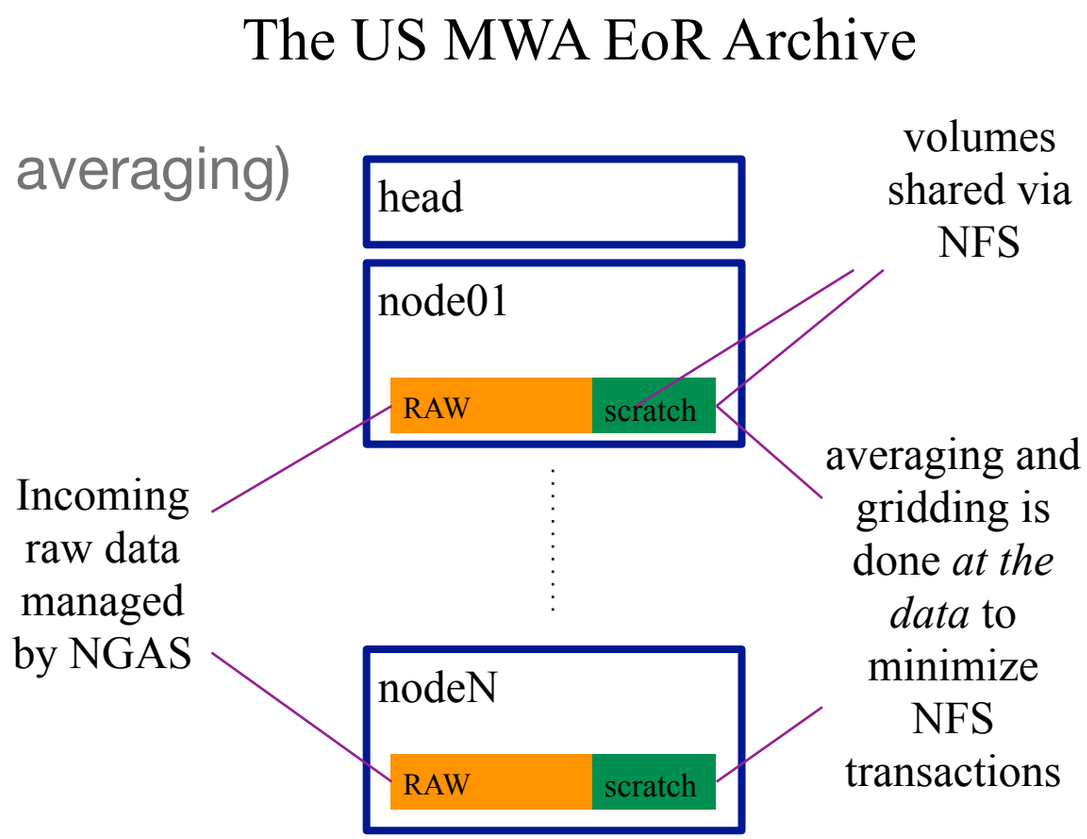

Computing on the data

- Shared resource with different usage models
  - transferring data
    (newly arriving data, hashing, results of averaging)
  - converting data
    (data bound, uses many cores)
  - subtract and average
    (data bound)
  - lst average
    (combines many data points, a corner turn)
  - other cpu intensive jobs (data archive is also largest and most capable machine)

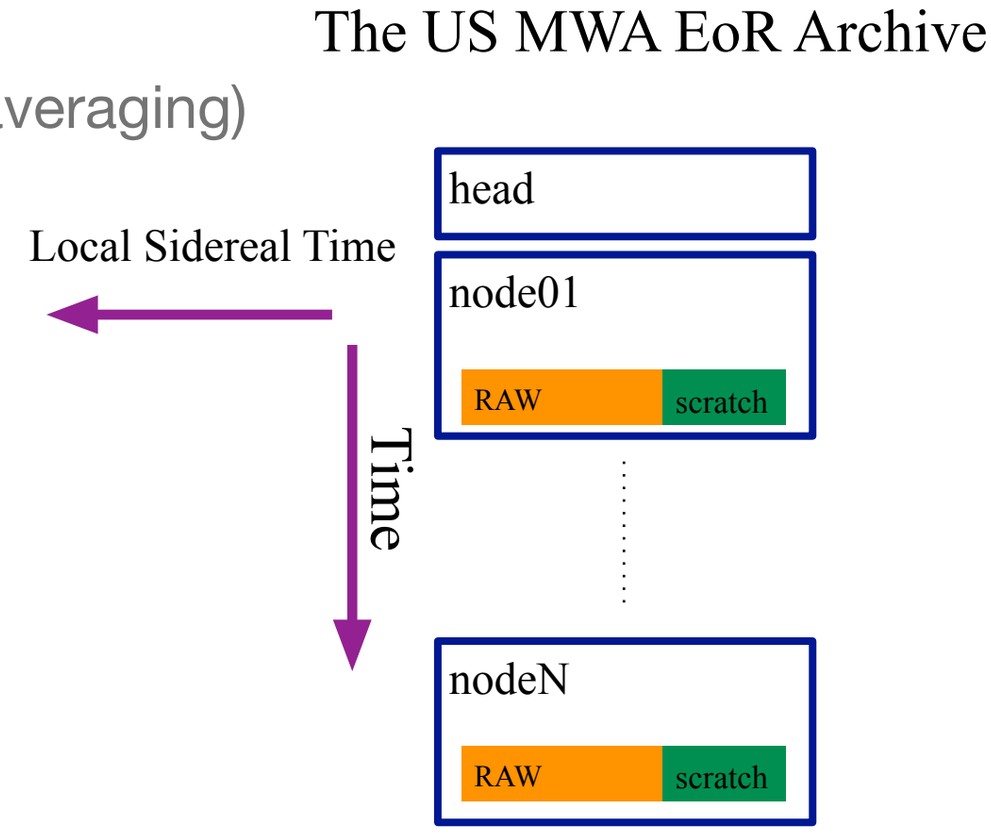

Data Processing Lessons

- transitioning to HPC is hard. most resources are not geared to data intensive applications. (100 $/TB on the open market, vs >350$/TB at a center)
    - nfs can hurt but this is what I can do in an afternoon
    - file formats rule the world
    - database connection limits are quickly reached, speed matters
    - managing data locations can get out of control quickly

Typical cluster problem

load on 8 core head node

```
top - 14:03:18 up 274 days, 21:52, 44 users,  load average: 40.30, 40.11, 40.0
Tasks: 1157 total,   2 running, 1153 sleeping,   0 stopped,   2 zombie
Cpu(s):  1.3%us, 27.6%sy,  0.1%ni, 68.1%id,  2.1%wa,  0.0%hi,  0.8%si,  0.0%st
Mem:  32850732k total, 18230012k used, 14620720k free,   372068k buffers
Swap: 16498680k total,   149336k used, 16349344k free,  4021216k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
  145 root      39  19     0    0    0 R 99.6  0.0 135426:41 kipmi0
19967 root      20   0 3714m 2.5g 2.4g S 86.8  7.8 113593:02 VBoxHeadless
10601 root      20   0 16964 1068  896 D 22.1  0.0   8439:39 mount.nfs
12627 root      20   0 5515m 4.1g 4.0g S  7.3 13.0 439292:13 VBoxHeadless
```

- management overhead step
  (from desktop to HPC)
  managing 1PB distributed on hardware with multiple use cases on <<99% hardware

Typical cluster problem

- mysteriously overloaded machine

load on 8 core head node

```
top - 14:03:18 up 274 days, 21:52, 44 users,  load average: 40.30, 40.11, 40.0
Tasks: 1157 total,   2 running, 1153 sleeping,   0 stopped,   2 zombie
Cpu(s):  1.3%us, 27.6%sy,  0.1%ni, 68.1%id,  2.1%wa,  0.0%hi,  0.8%si,  0.0%st
Mem:  32850732k total, 18230012k used, 14620720k free,   372068k buffers
Swap: 16498680k total,   149336k used, 16349344k free,  4021216k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
  145 root      39  19     0    0    0 R 99.6  0.0 135426:41 kipmi0
19967 root      20   0 3714m 2.5g 2.4g S 86.8  7.8 113593:02 VBoxHeadless
10601 root      20   0 16964 1068  896 D 22.1  0.0   8439:39 mount.nfs
12627 root      20   0 5515m 4.1g 4.0g S  7.3 13.0 439292:13 VBoxHeadless
```

- management overhead step
  (from desktop to HPC)
  managing 1PB distributed on hardware with multiple use cases on <<99% hardware

Typical cluster problem

- mysteriously overloaded machine

load on 8 core head node

```
top - 14:03:18 up 274 days, 21:52, 44 users,  load average: 40.30, 40.11, 40.0
Tasks: 1157 total,   2 running, 1153 sleeping,   0 stopped,   2 zombie
Cpu(s):  1.3%us, 27.6%sy,  0.1%ni, 68.1%id,  2.1%wa,  0.0%hi,  0.8%si,  0.0%st
Mem:  32850732k total, 18230012k used, 14620720k free,   372068k buffers
Swap: 16498680k total,   149336k used, 16349344k free,  4021216k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
  145 root      39  19     0    0    0 R 99.6  0.0 135426:41 kipmi0
19967 root      20   0 3714m 2.5g 2.4g S 86.8  7.8 113593:02 VBoxHeadless
10601 root      20   0 16964 1068  896 D 22.1  0.0   8439:39 mount.nfs
12627 root      20   0 5515m 4.1g 4.0g S  7.3 13.0 439292:13 VBoxHeadless
```

- management overhead step
  (from desktop to HPC)
  managing 1PB distributed on hardware with multiple use cases on <<99% hardware

Typical cluster problem

- mysteriously overloaded machine
- still a mystery as we speak but definitely io related

load on 8 core head node

```
top - 14:03:18 up 274 days, 21:52, 44 users,  load average: 40.30, 40.11, 40.0
Tasks: 1157 total,   2 running, 1153 sleeping,   0 stopped,   2 zombie
Cpu(s):  1.3%us, 27.6%sy,  0.1%ni, 68.1%id,  2.1%wa,  0.0%hi,  0.8%si,  0.0%st
Mem:  32850732k total, 18230012k used, 14620720k free,   372068k buffers
Swap: 16498680k total,   149336k used, 16349344k free,  4021216k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
  145 root      39  19     0    0    0 R 99.6  0.0 135426:41 kipmi0
19967 root      20   0 3714m 2.5g 2.4g S 86.8  7.8 113593:02 VBoxHeadless
10601 root      20   0 16964 1068  896 D 22.1  0.0   8439:39 mount.nfs
12627 root      20   0 5515m 4.1g 4.0g S  7.3 13.0 439292:13 VBoxHeadless
```

- management overhead step
  (from desktop to HPC)
  managing 1PB distributed on hardware with multiple use cases on <<99% hardware

Conclusions: Challenges

Conclusions: Challenges

Being smart about metadata

Conclusions: Challenges

Being smart about metadata

Tracking Data through its lifetime

Conclusions: Challenges

Being smart about metadata

Tracking Data through its lifetime

Bootstrapping onto HPC