The I/O Challenges of Ultrascale Visualization for the Square Kilometre Array and its Pre-cursers

Andreas Wicenec
International Centre for Radio Astronomy Research
Perth, Western Australia

ICRAR International Centre for Radio Astronomy Research

Intro

"The output from leading-edge scientific simulations is so voluminous and complex that advanced visualization techniques are necessary to interpret the calculated results."

This talk is about a few upcoming and planned astronomical instruments producing multi-dimensional data sets at stunning rates and volumes using HPC as an integrated part of the data flow.

Eyes on the Sky

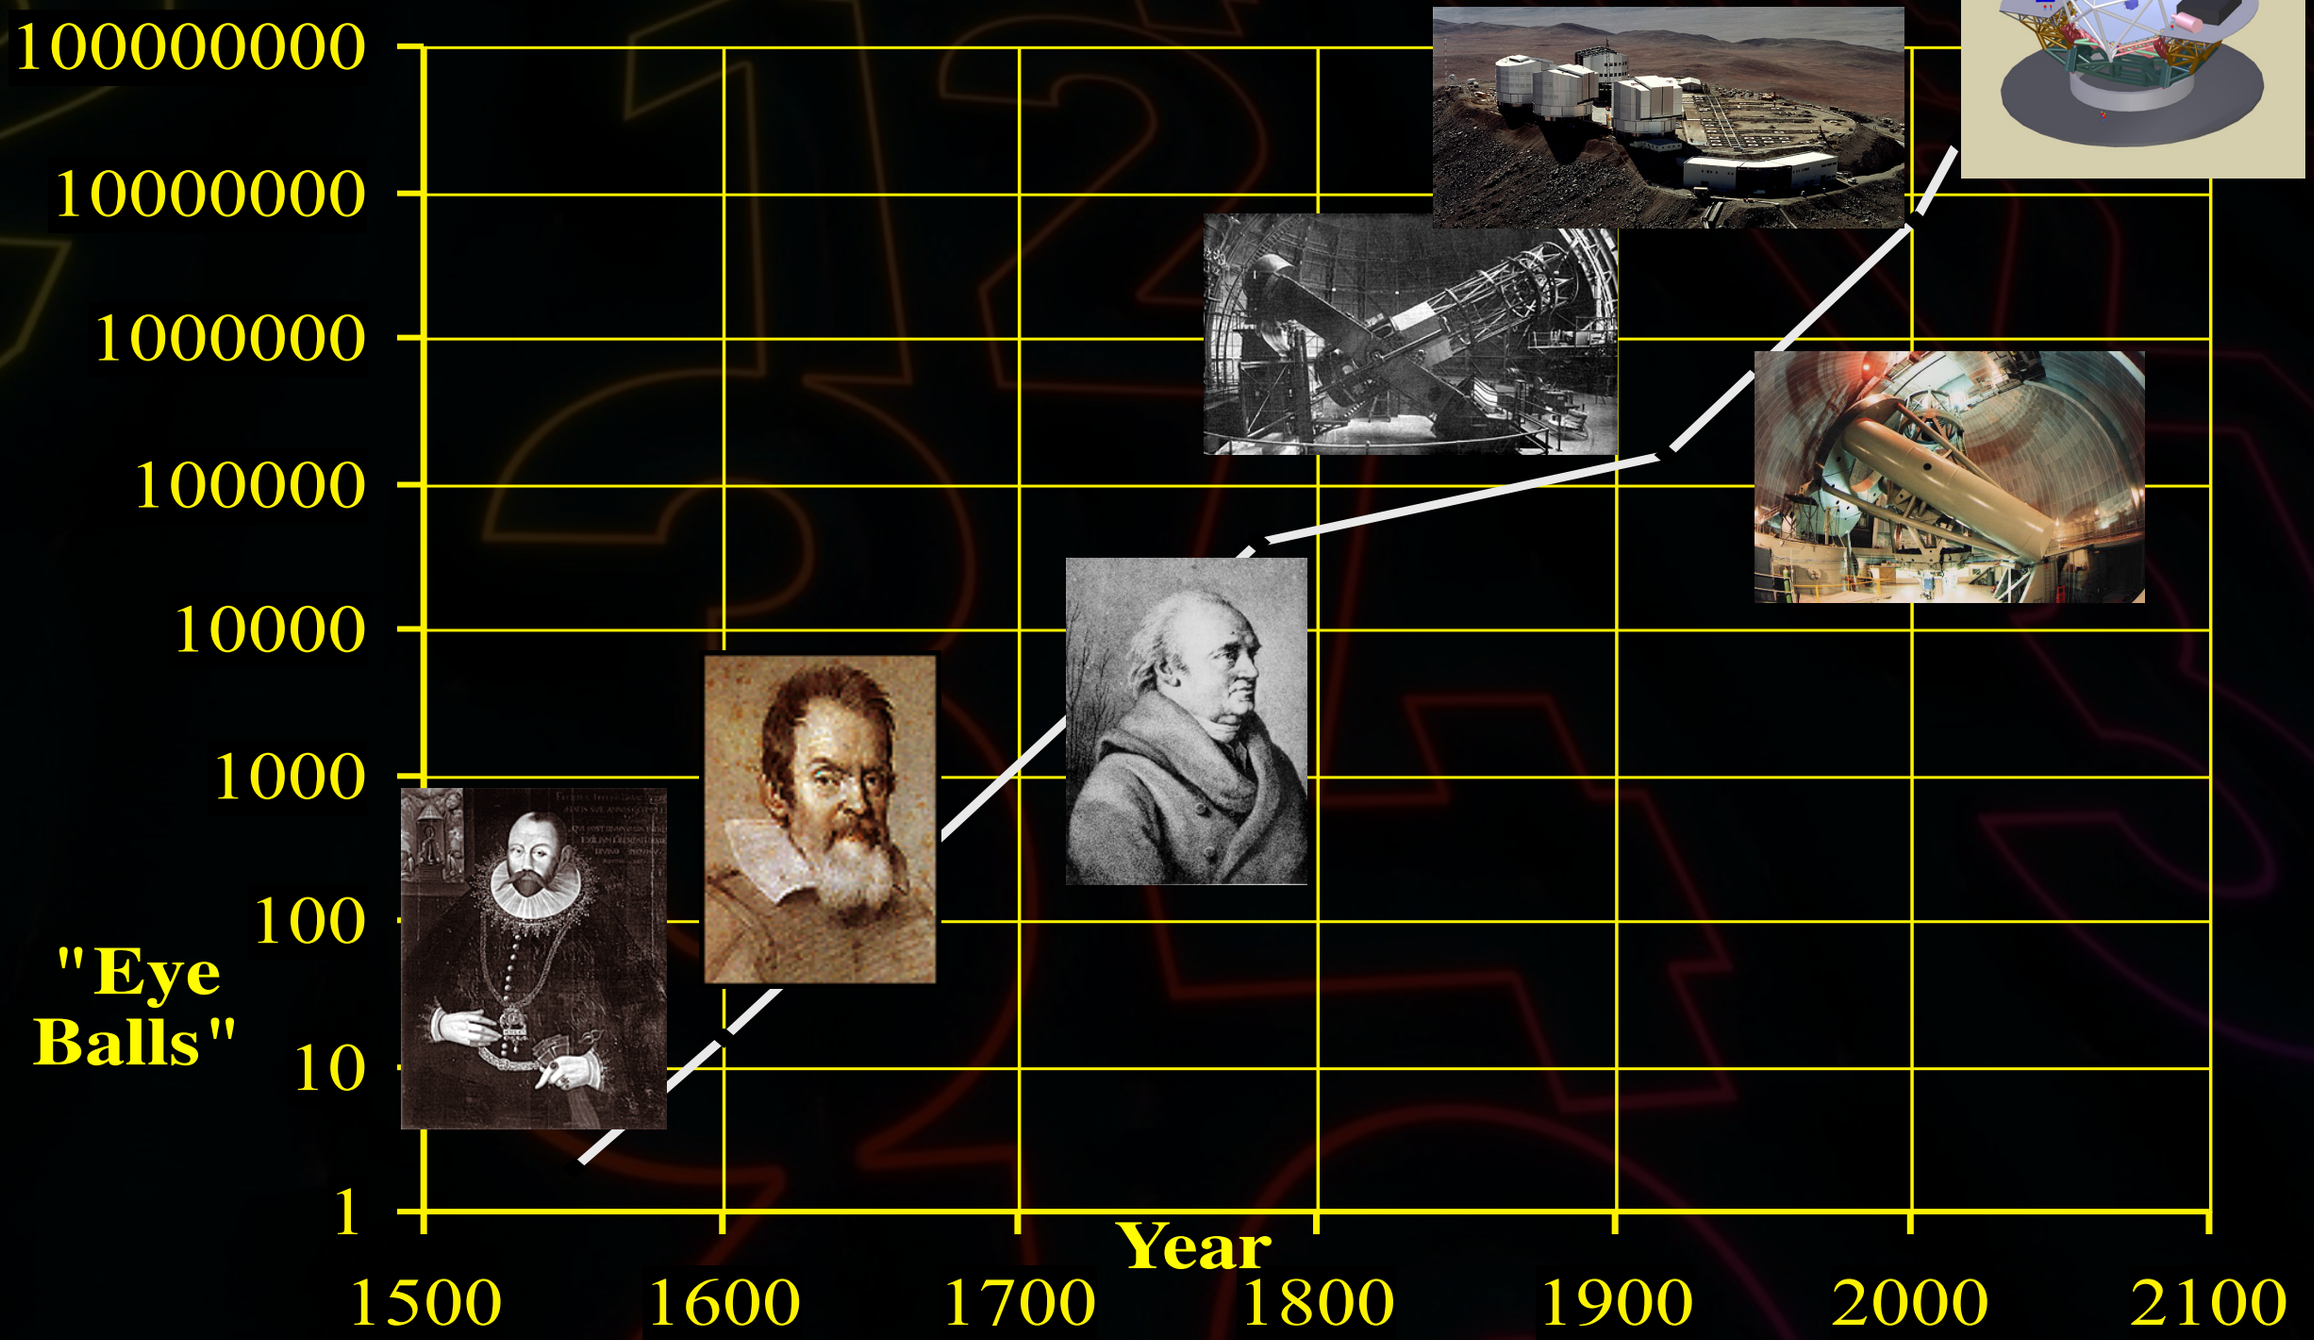

Eyes on the Sky

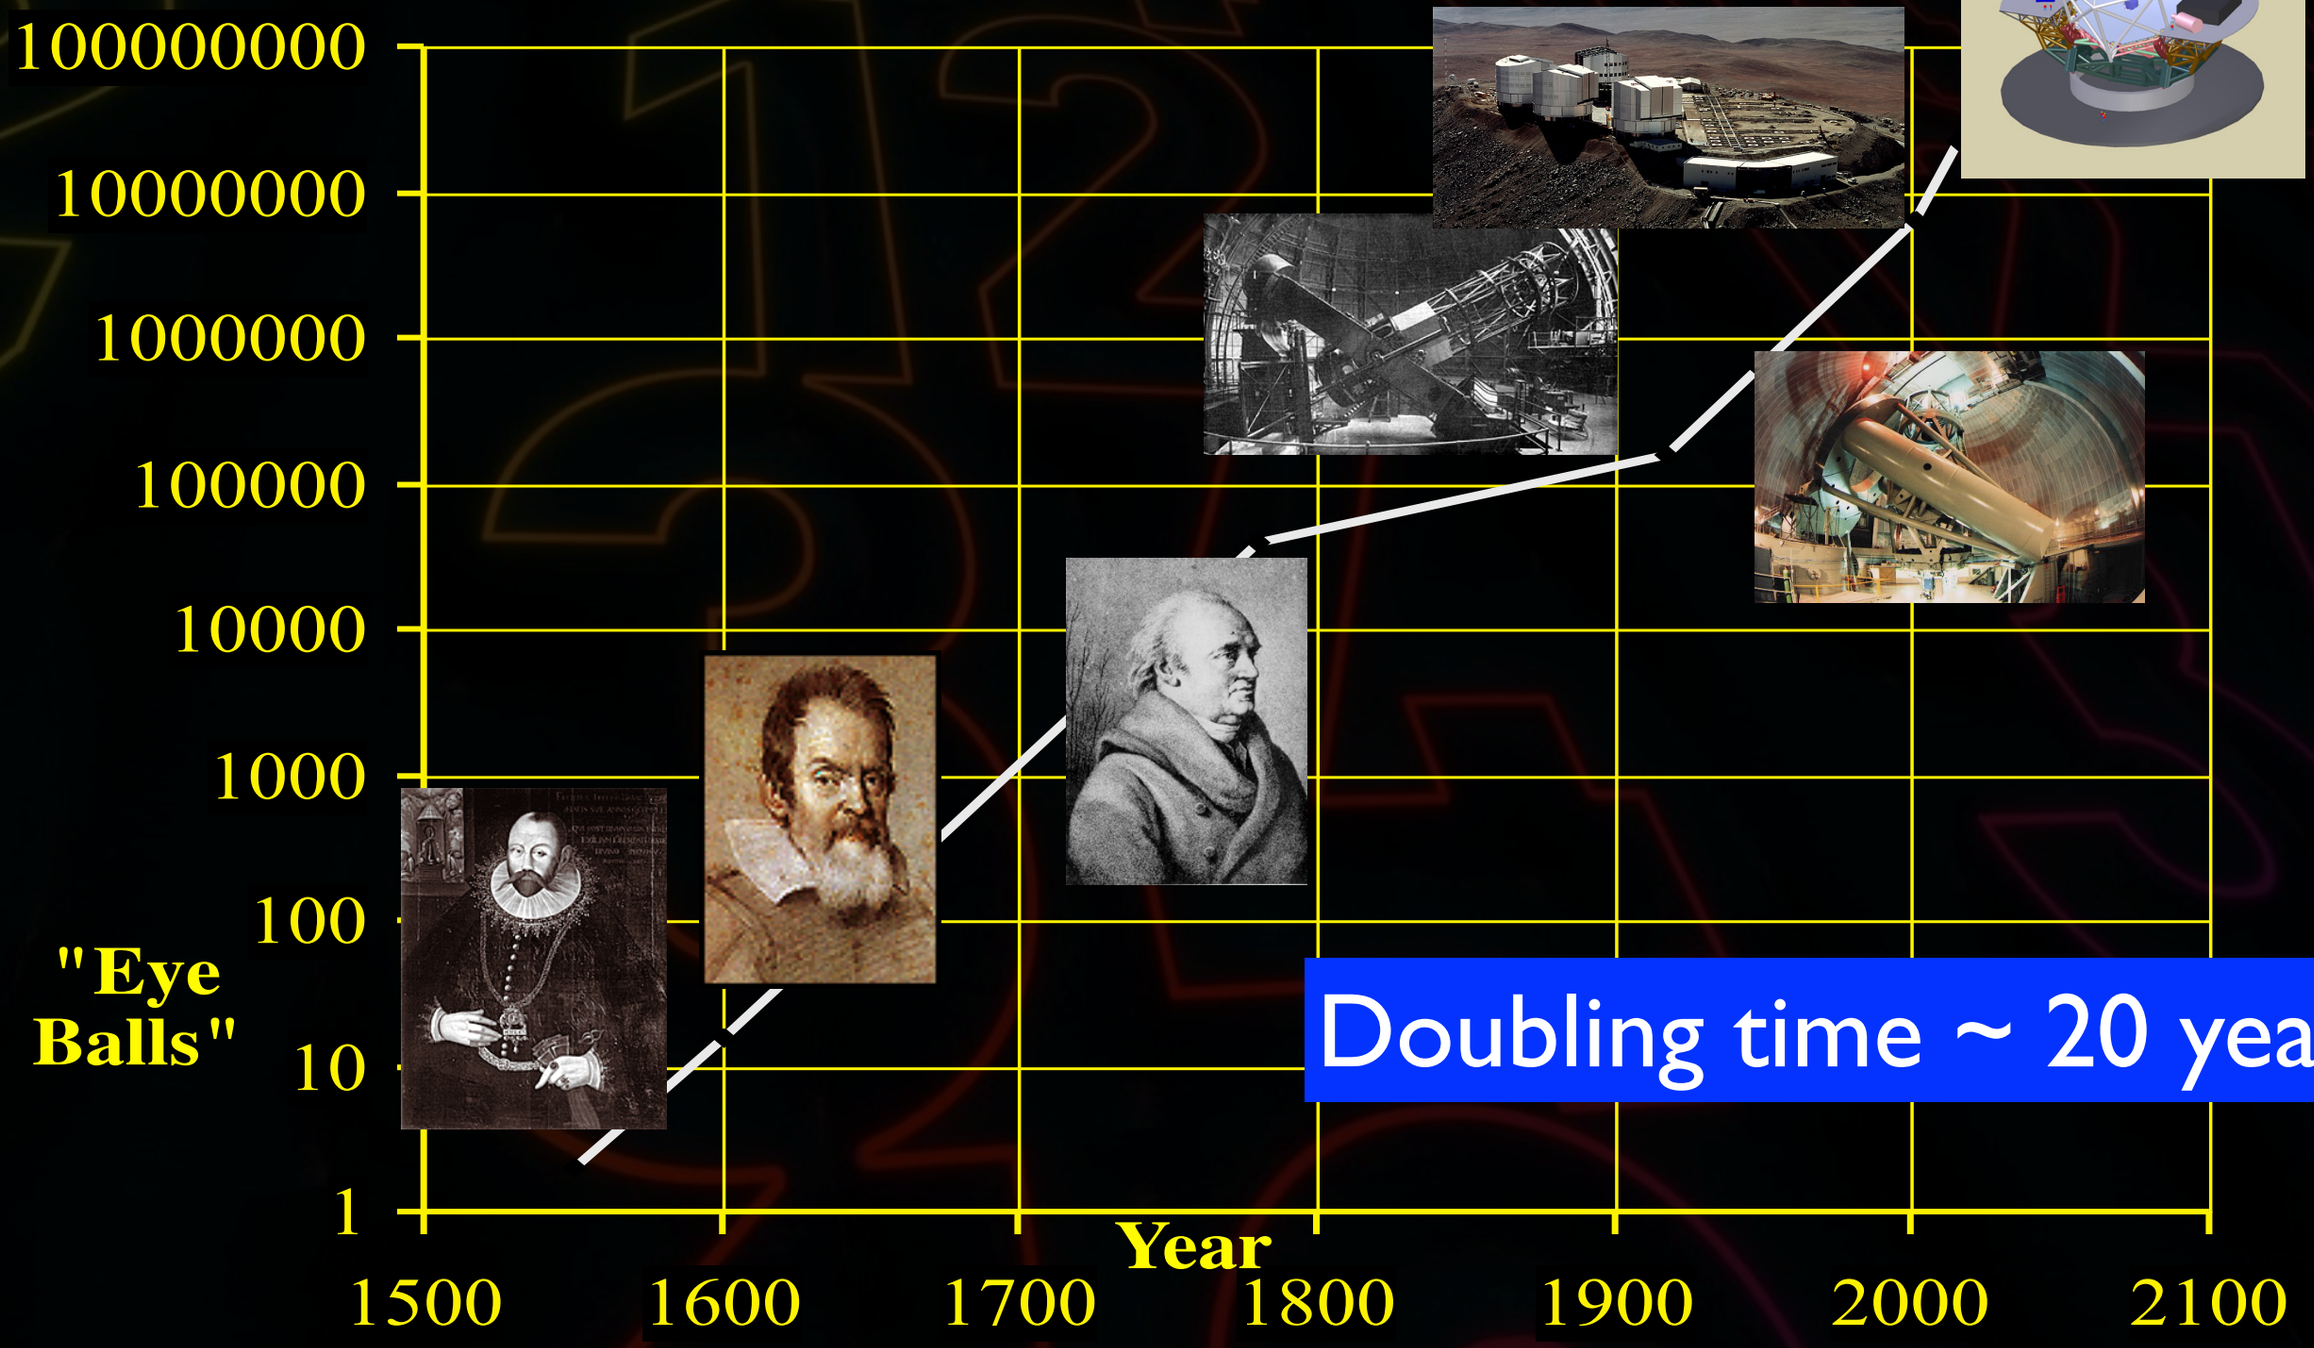

Gathering numbers

Gathering numbers

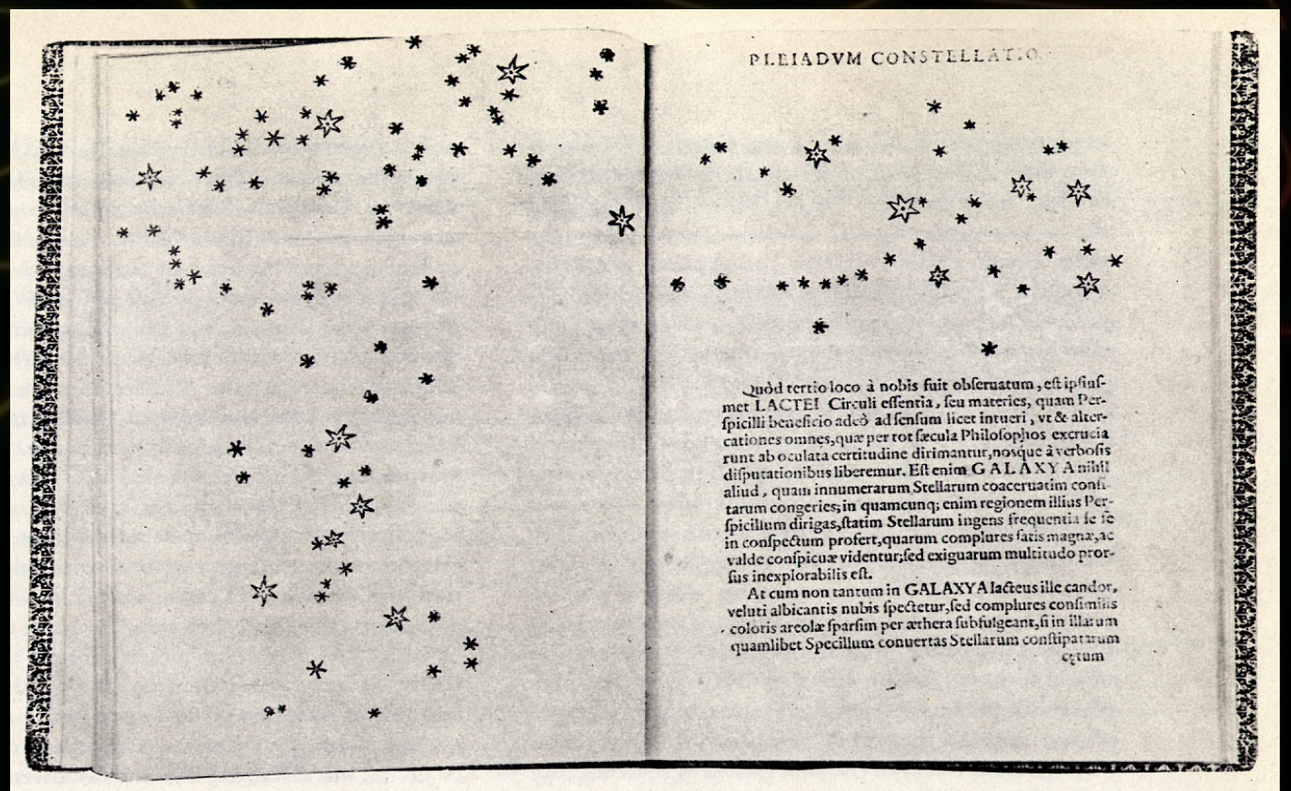

1610
Nearby stars

Gathering numbers

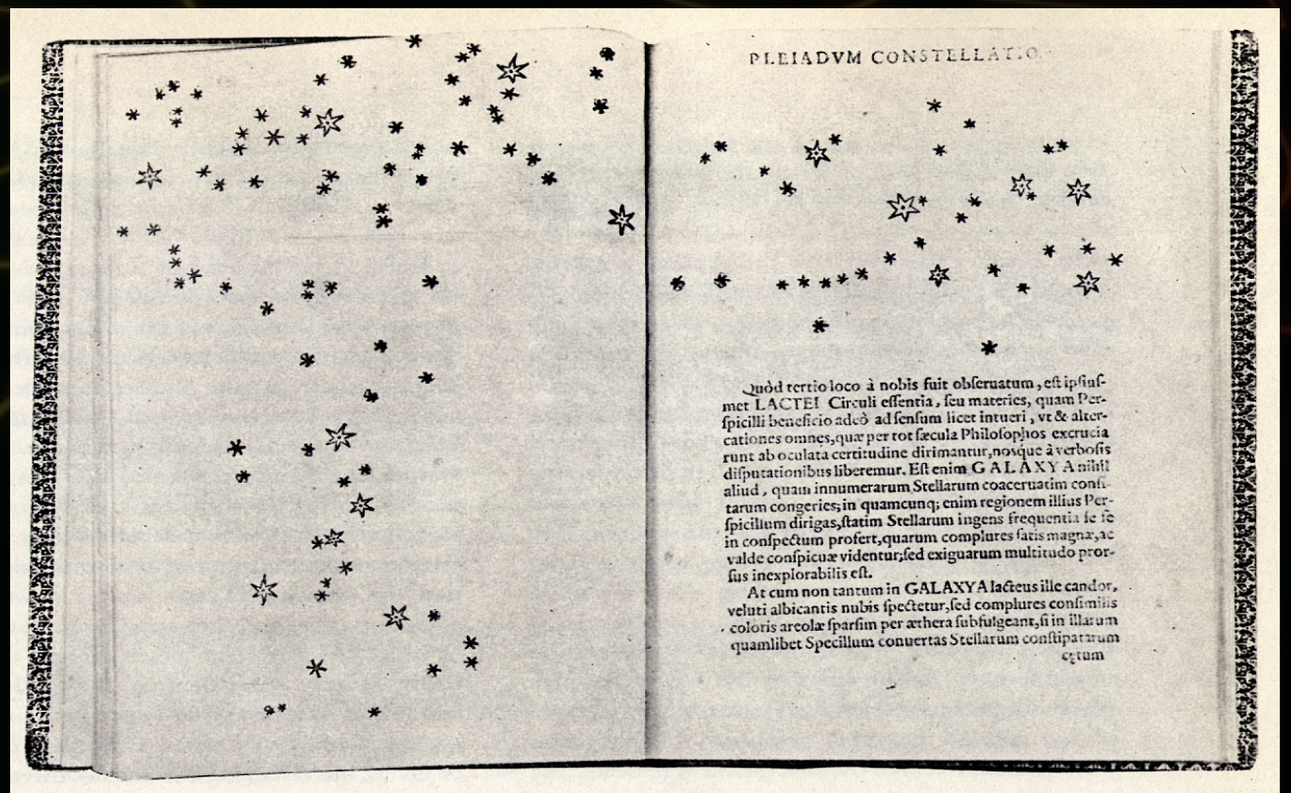
PLEIADVM CONSTELLATIO.

Quòd tertio loco à nobis fuit obſeruatum, eſt ipſiuſmet LACTEI Circuli eſſentia, ſeu materies, quam Perſpicilli beneficio adeò ad ſenſum licet intueri, vt & altercationes omnes, quæ per tot ſæcula Philoſophos excruciarunt ab oculata certitudine dirimantur, nosque à verboſis diſputationibus liberemur. Eſt enim GALAXYA nihil aliud, quam innumerarum Stellarum coaceruatim conſitarum congeries; in quamcunq; enim regionem illius Perſpicillum dirigas, ſtatim Stellarum ingens frequentia ſe ſe in conſpectum profert, quarum complures ſatis magnæ, ac valde conſpicuæ videntur; ſed exiguarum multitudo prorſus inexplorabilis eſt.

At cum non tantum in GALAXYA lacteus ille candor, veluti albicantis nubis ſpectetur, ſed complures conſimilis coloris arcolæ ſparſim per æthera ſubfulgeant, ſi in illarum quamlibet Specillum conuertas Stellarum conſipatarum

etum

1610
Nearby stars

1845
Ink sketch of nearby galaxy

Gathering numbers

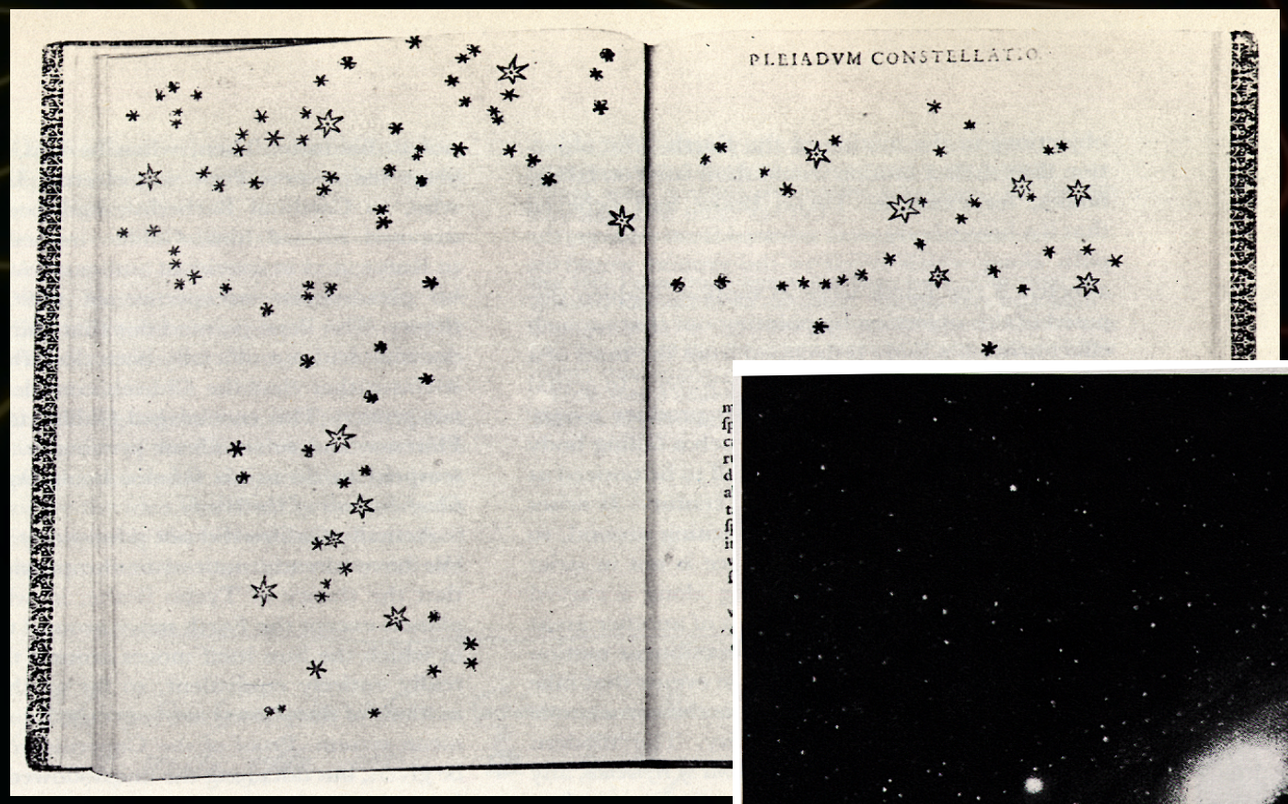

1610
Nearby stars

1880
First photographs

1845
Ink sketch of nearby galaxy

Numbers per night

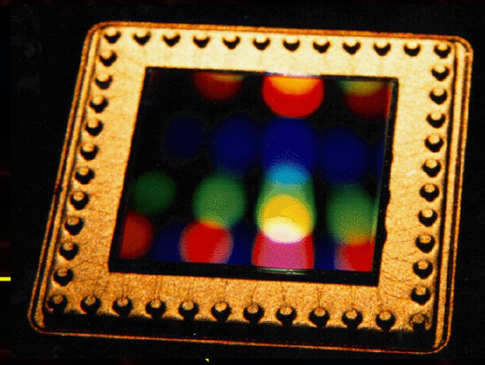

1E+12
1E+11
1E+10
1E+09
100000000
10000000
1000000
100000
10000
1000
100

1500 1600 1700 1800 1900 2000

Numbers/ night

Year

Numbers per night

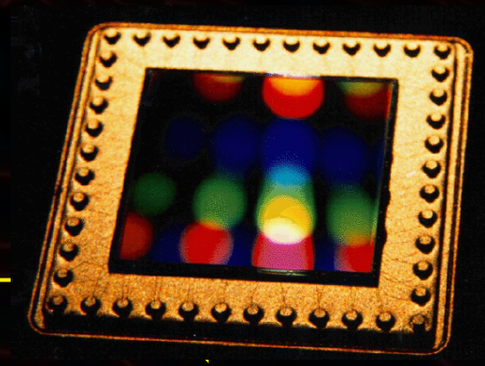

1E+12
1E+11
1E+10
1E+09
100000000
10000000
1000000
100000
10000
1000
100

1500 1600 1700 1800 1900 2000

Numbers/ night

Year

Doubling time < 1 year

The deluge continues

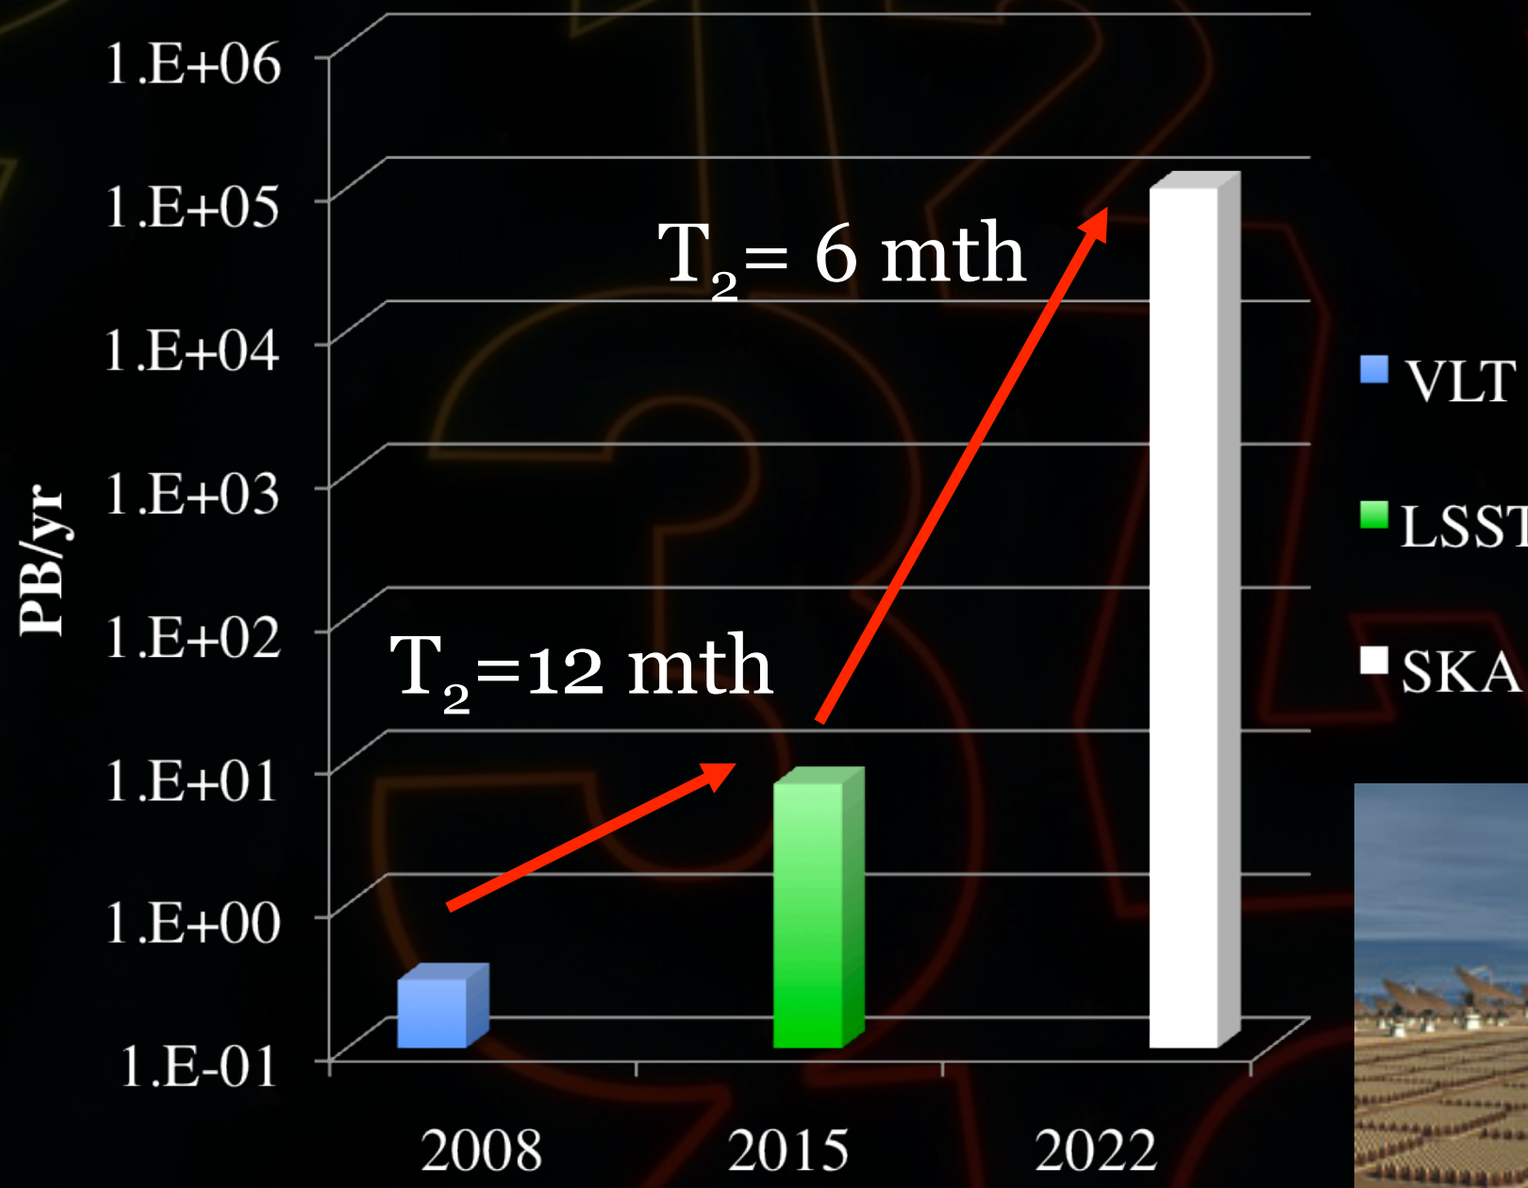

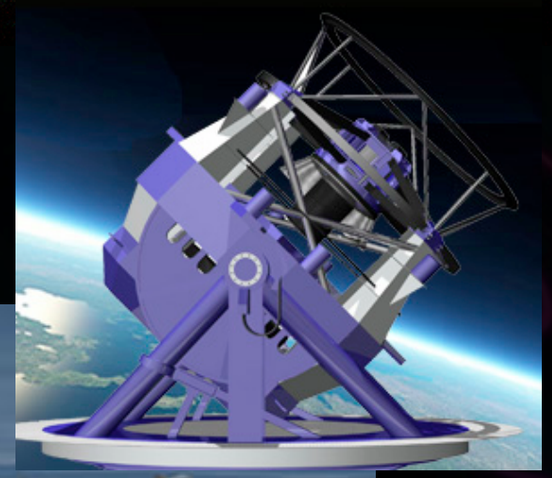

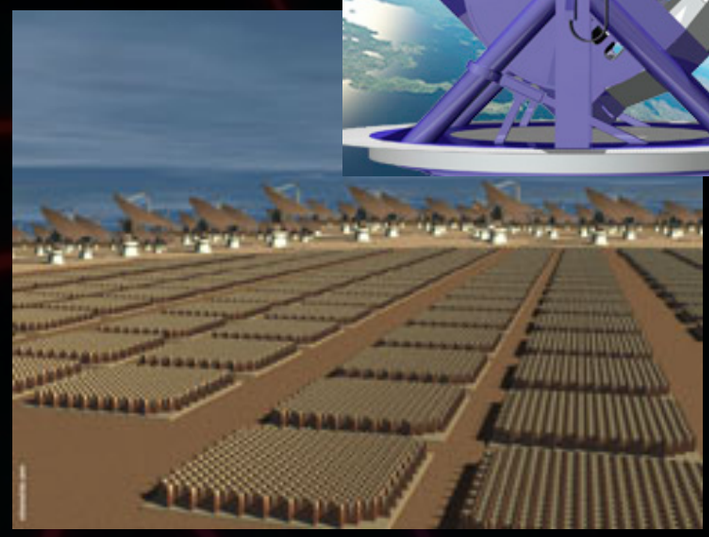

The deluge continues

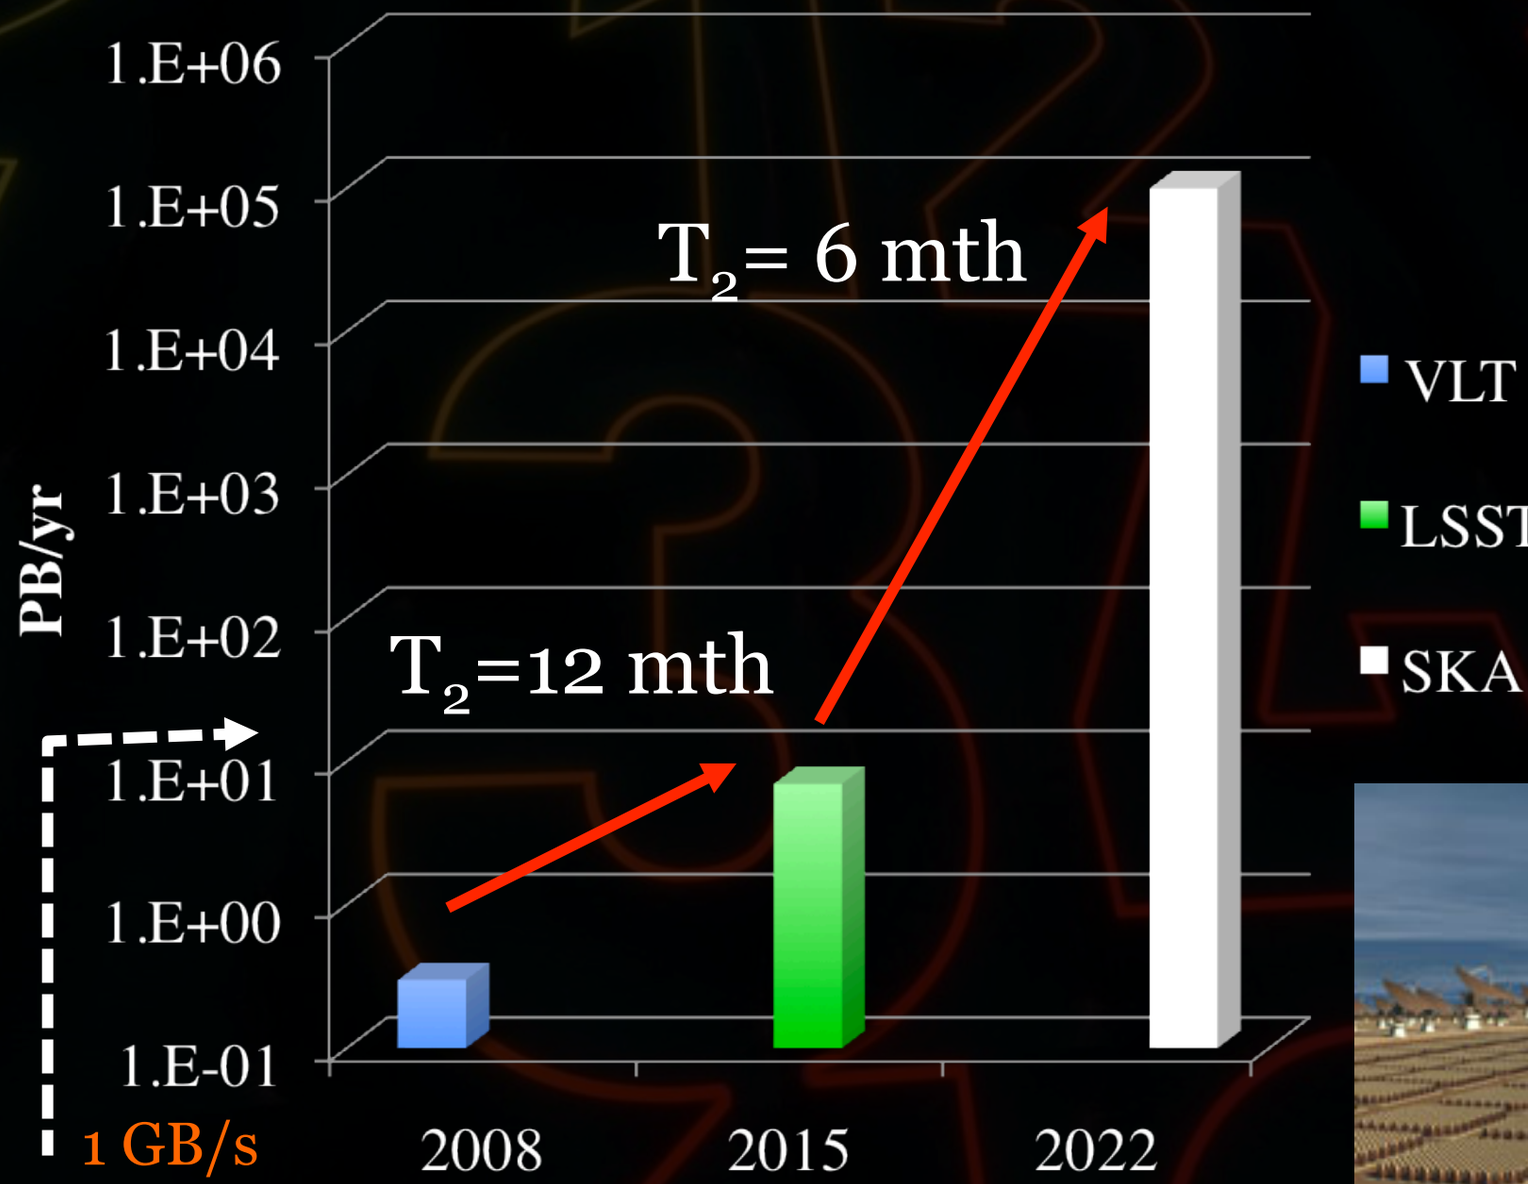

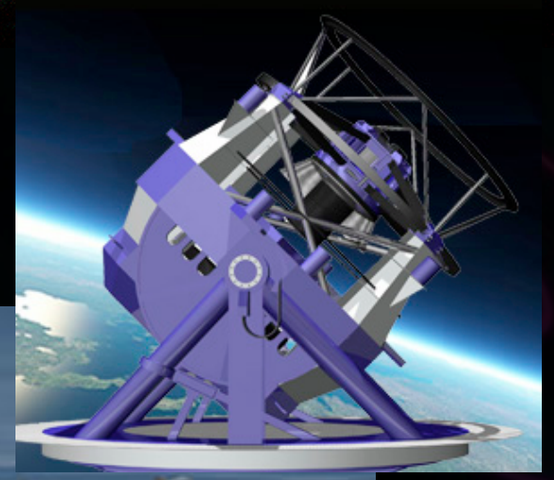

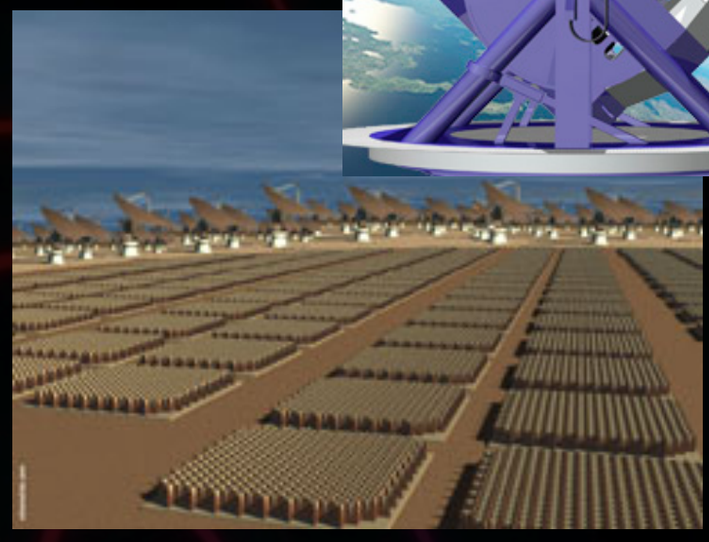

The deluge continues

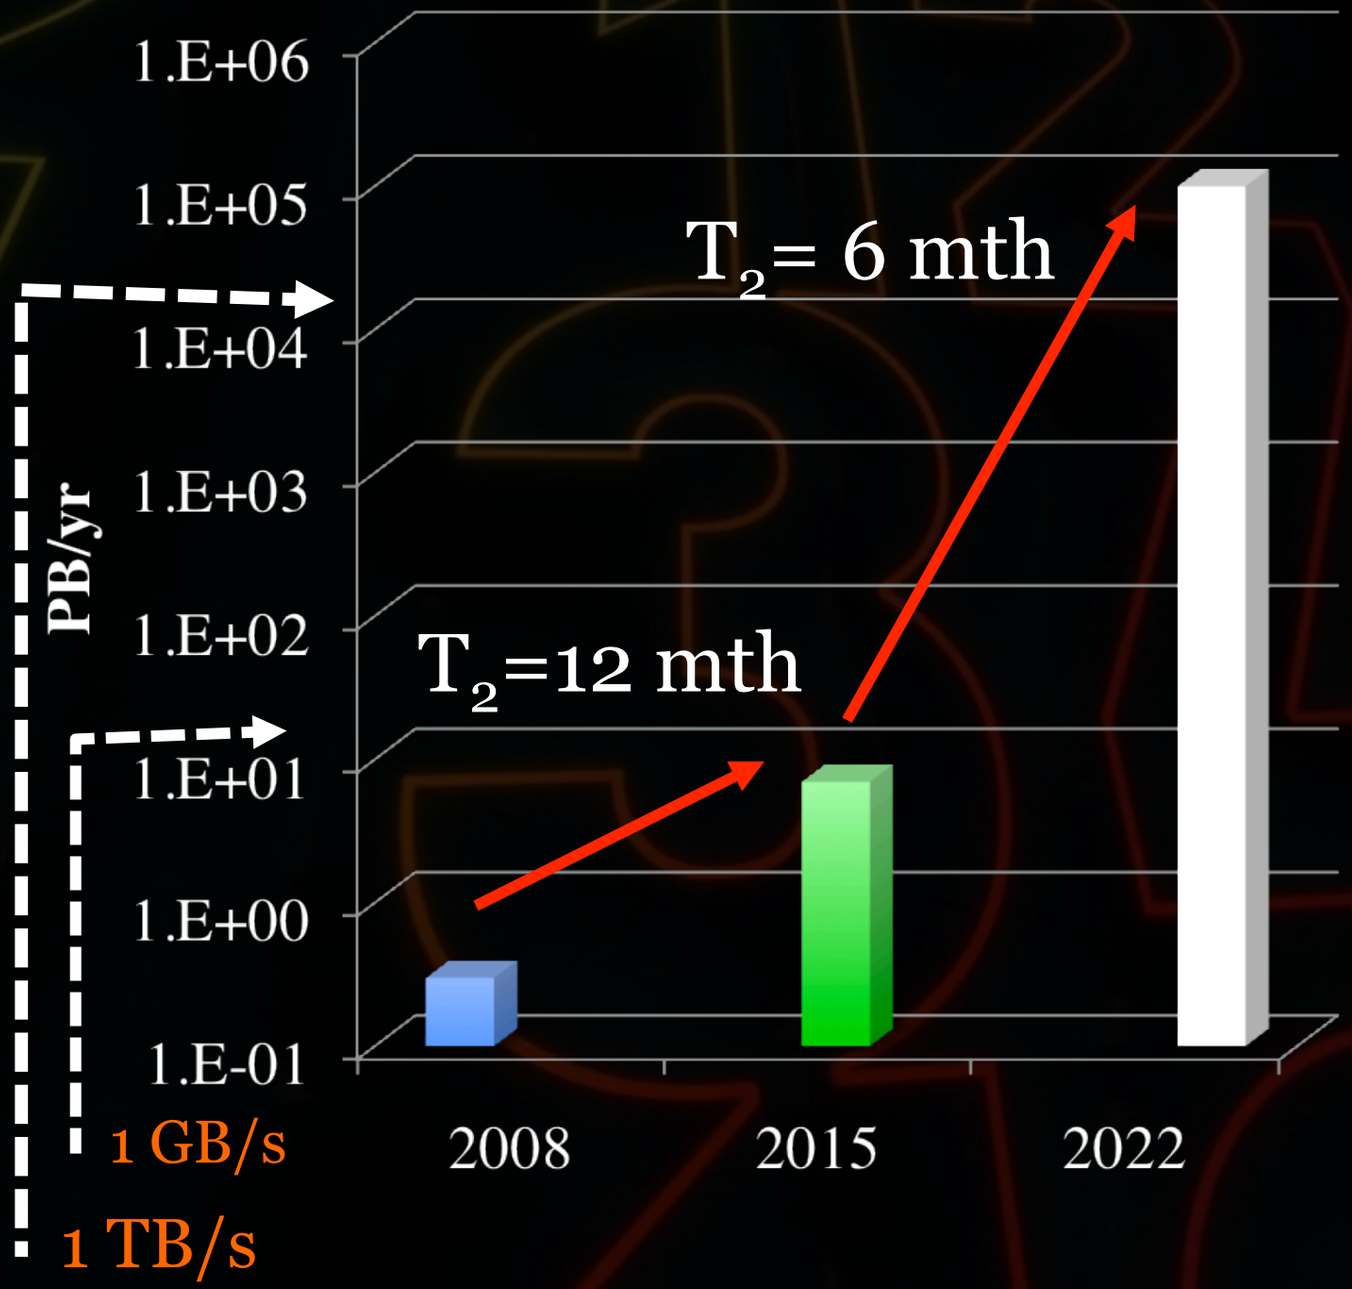

VLT

LSST

SKA

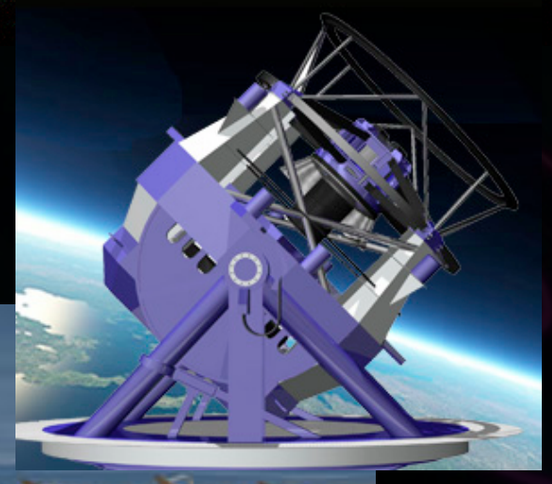

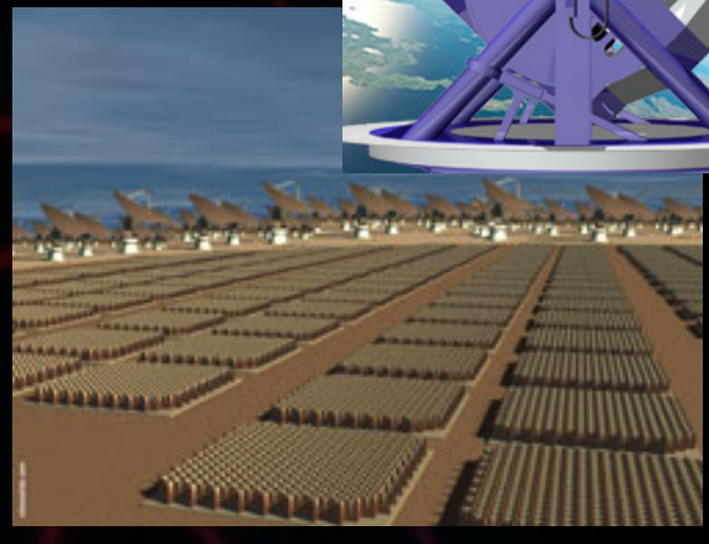

The deluge continues

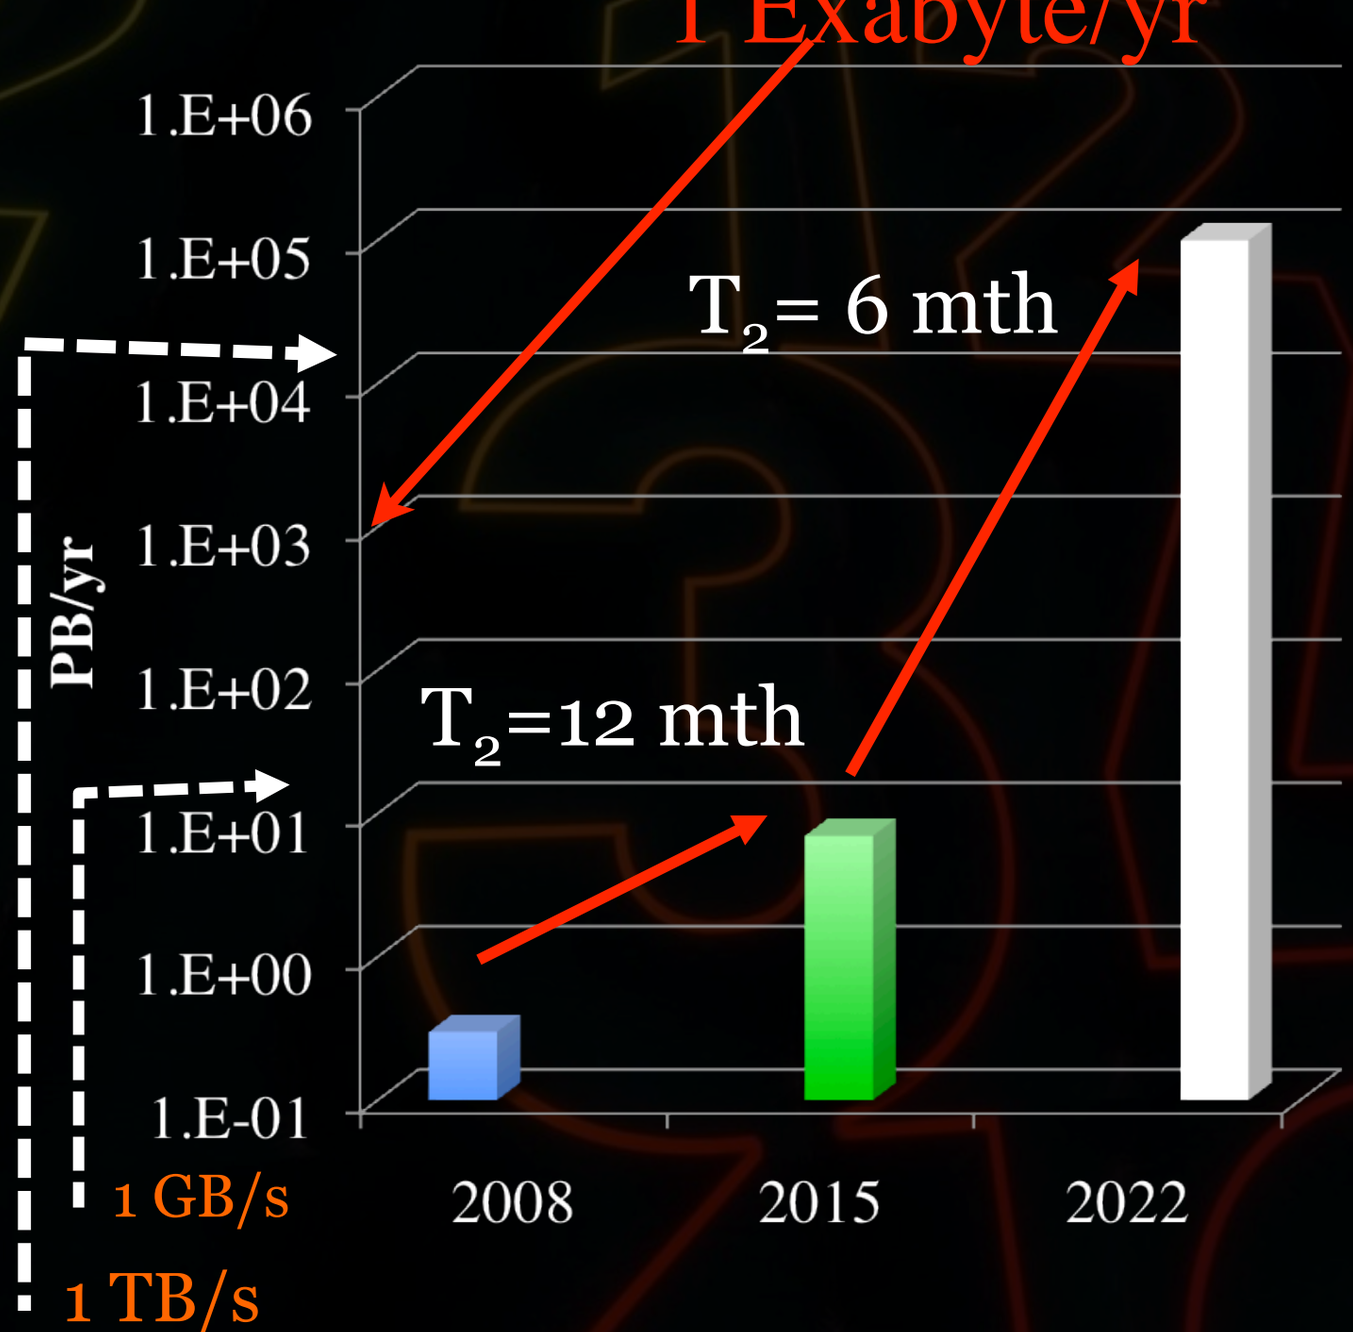

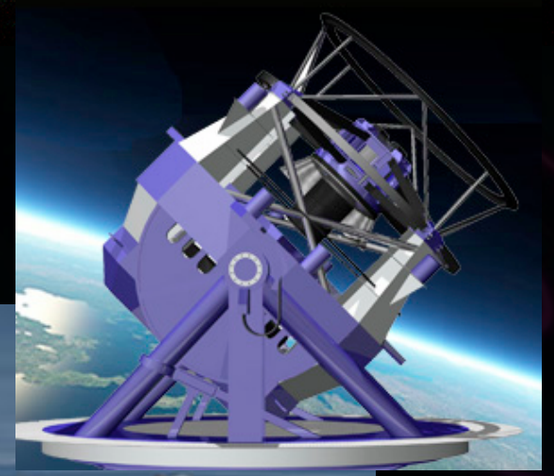

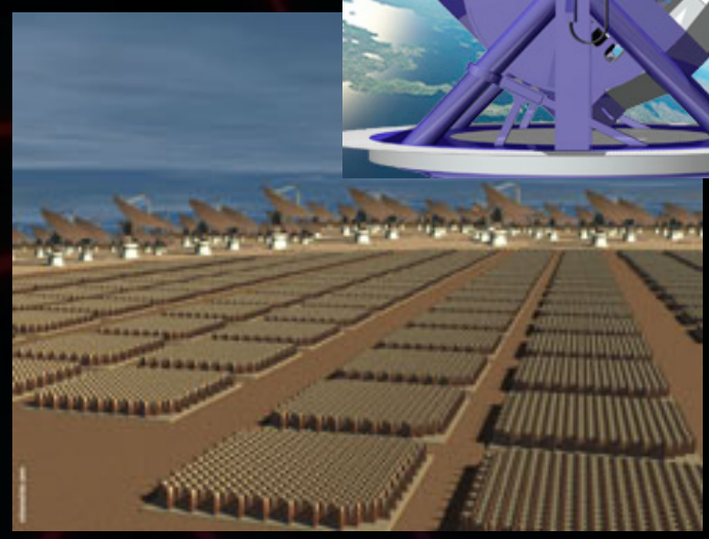

The SKA

The Square Kilometre Array (SKA) will be the largest international astronomical facility of the 21st century.

It will consist of up to 3000 dishes and hundreds of aperture arrays distributed over a range of up to 5000 km. The total collecting area will be of the order of one square kilometre.

It will observe the sky in radio frequencies between 50 MHz and 35 GHz.

The main science goals are in the area of the very early universe.

In 2009 the world produced 1,000,000,000,000,000,000 bytes of information. The SKA could potentially produce this data volume in one day.

ASKAP & MeerKAT

The Australian SKA Pathfinder (ASKAP) and the South African MeerKAT are currently under construction and represent 1% SKA each.

Technology testbeds and scientific facilities.

Will produce science data cubes of about 6 TB each. ASKAP is a wide field survey instrument and will produce several thousand cubes per survey. 10 surveys have been proposed.

Crunching the numbers

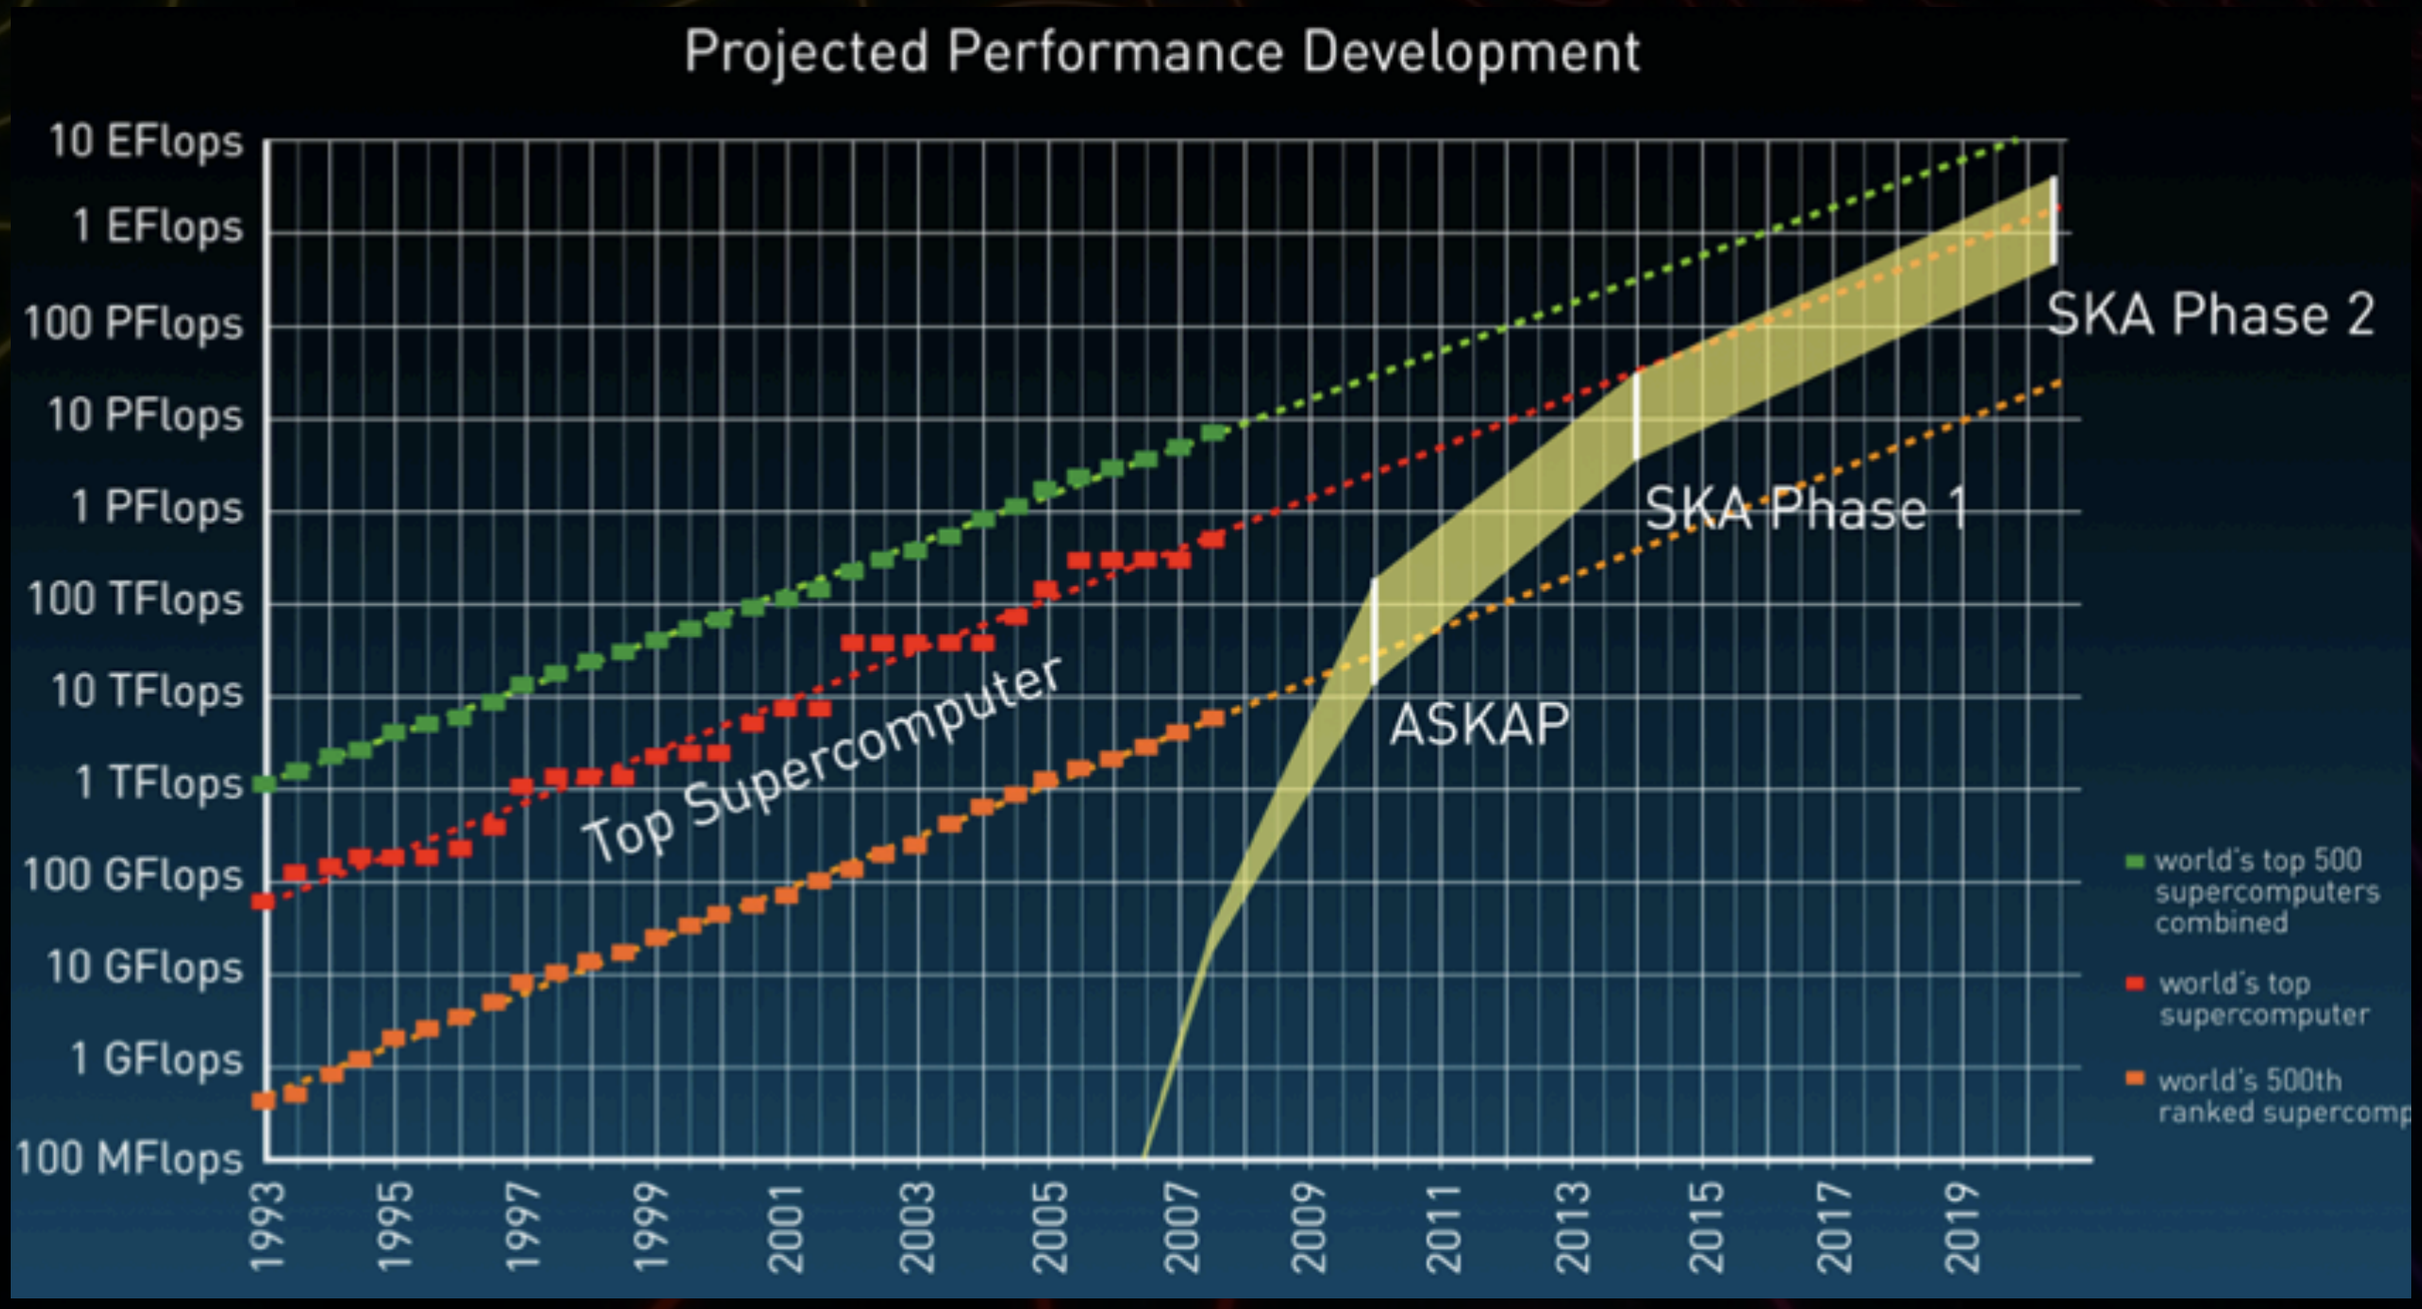

credit: T. Cornwell

Reading numbers

- Data comes in cubes
- SKA Pathfinder Cubes ~ 6 TB
  which implies 600 sec read time at 10GB/sec
- typical survey consists of ~1500 cubes = 10 days read time
- would like 100-1000 GB/sec for on-demand processing single cubes and cube groups.

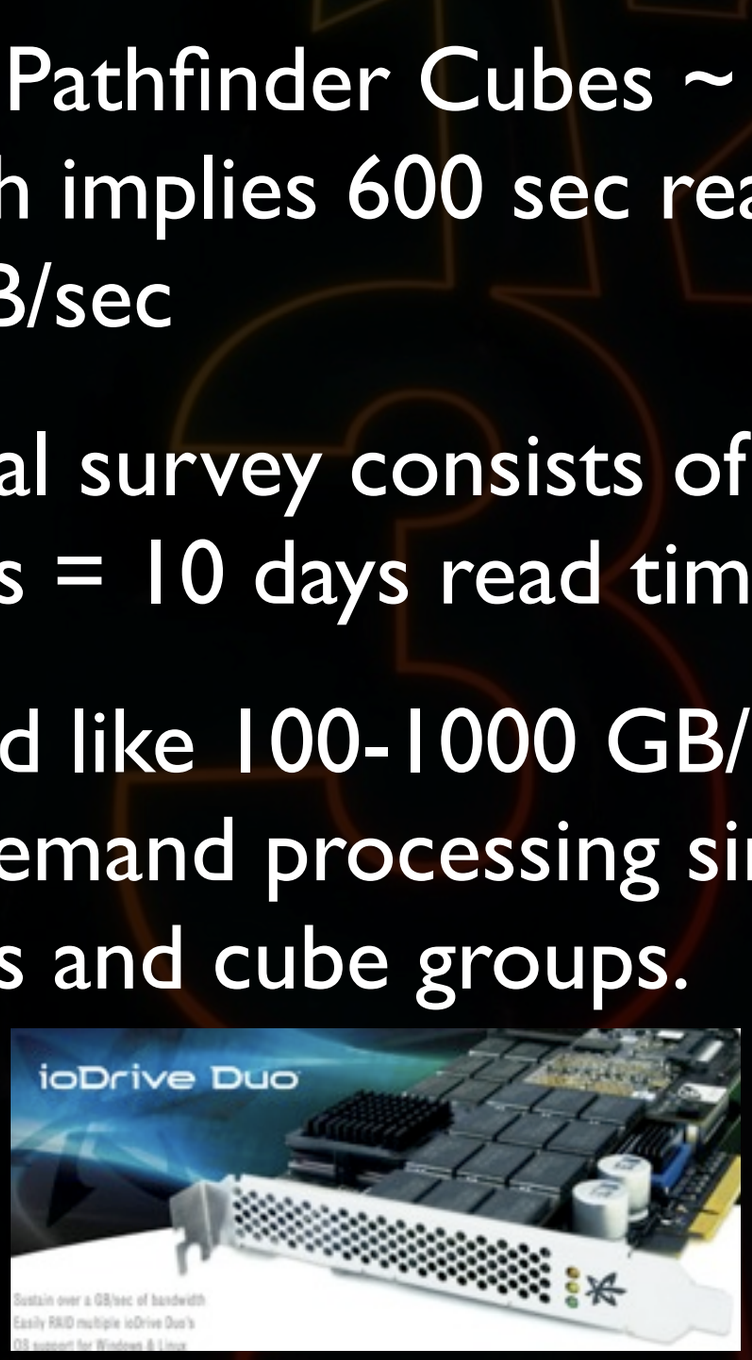

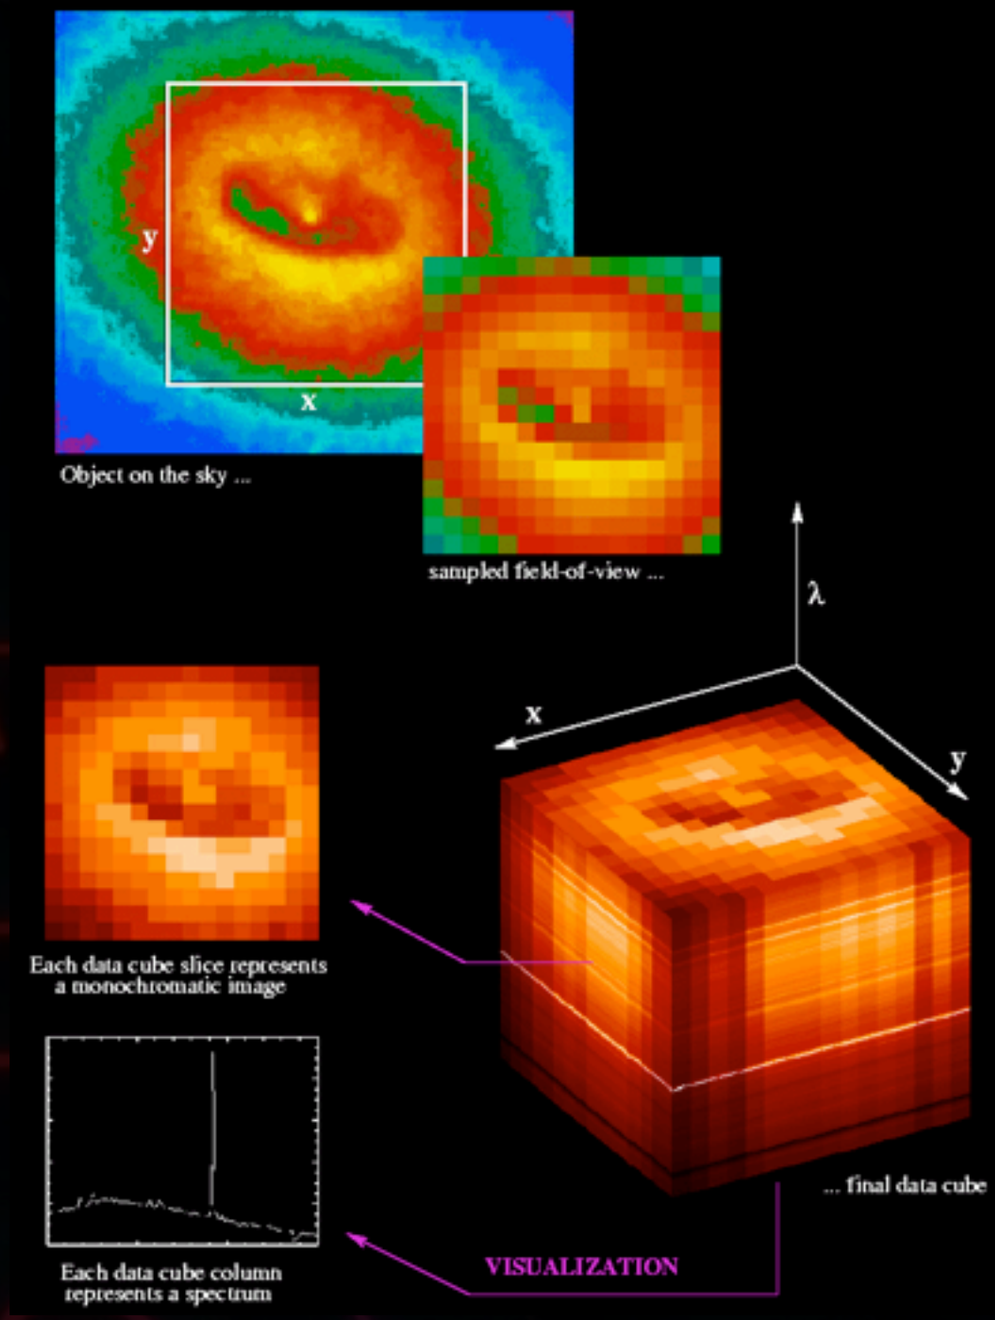

New Type of SC

- I/O is THE bottleneck for this kind of science
- Joining of HPC, highest performance storage and database technology == Data Machine
- Dedicated HPC design for Data Intensive Research (DIR) and visualisation
- Integration and optimisation of job scheduling and data movement from lower to highest tiers required.
- Integration of data movement from high performance storage to host and device memory.

DIRP Data Flow

Visualisation Use Case

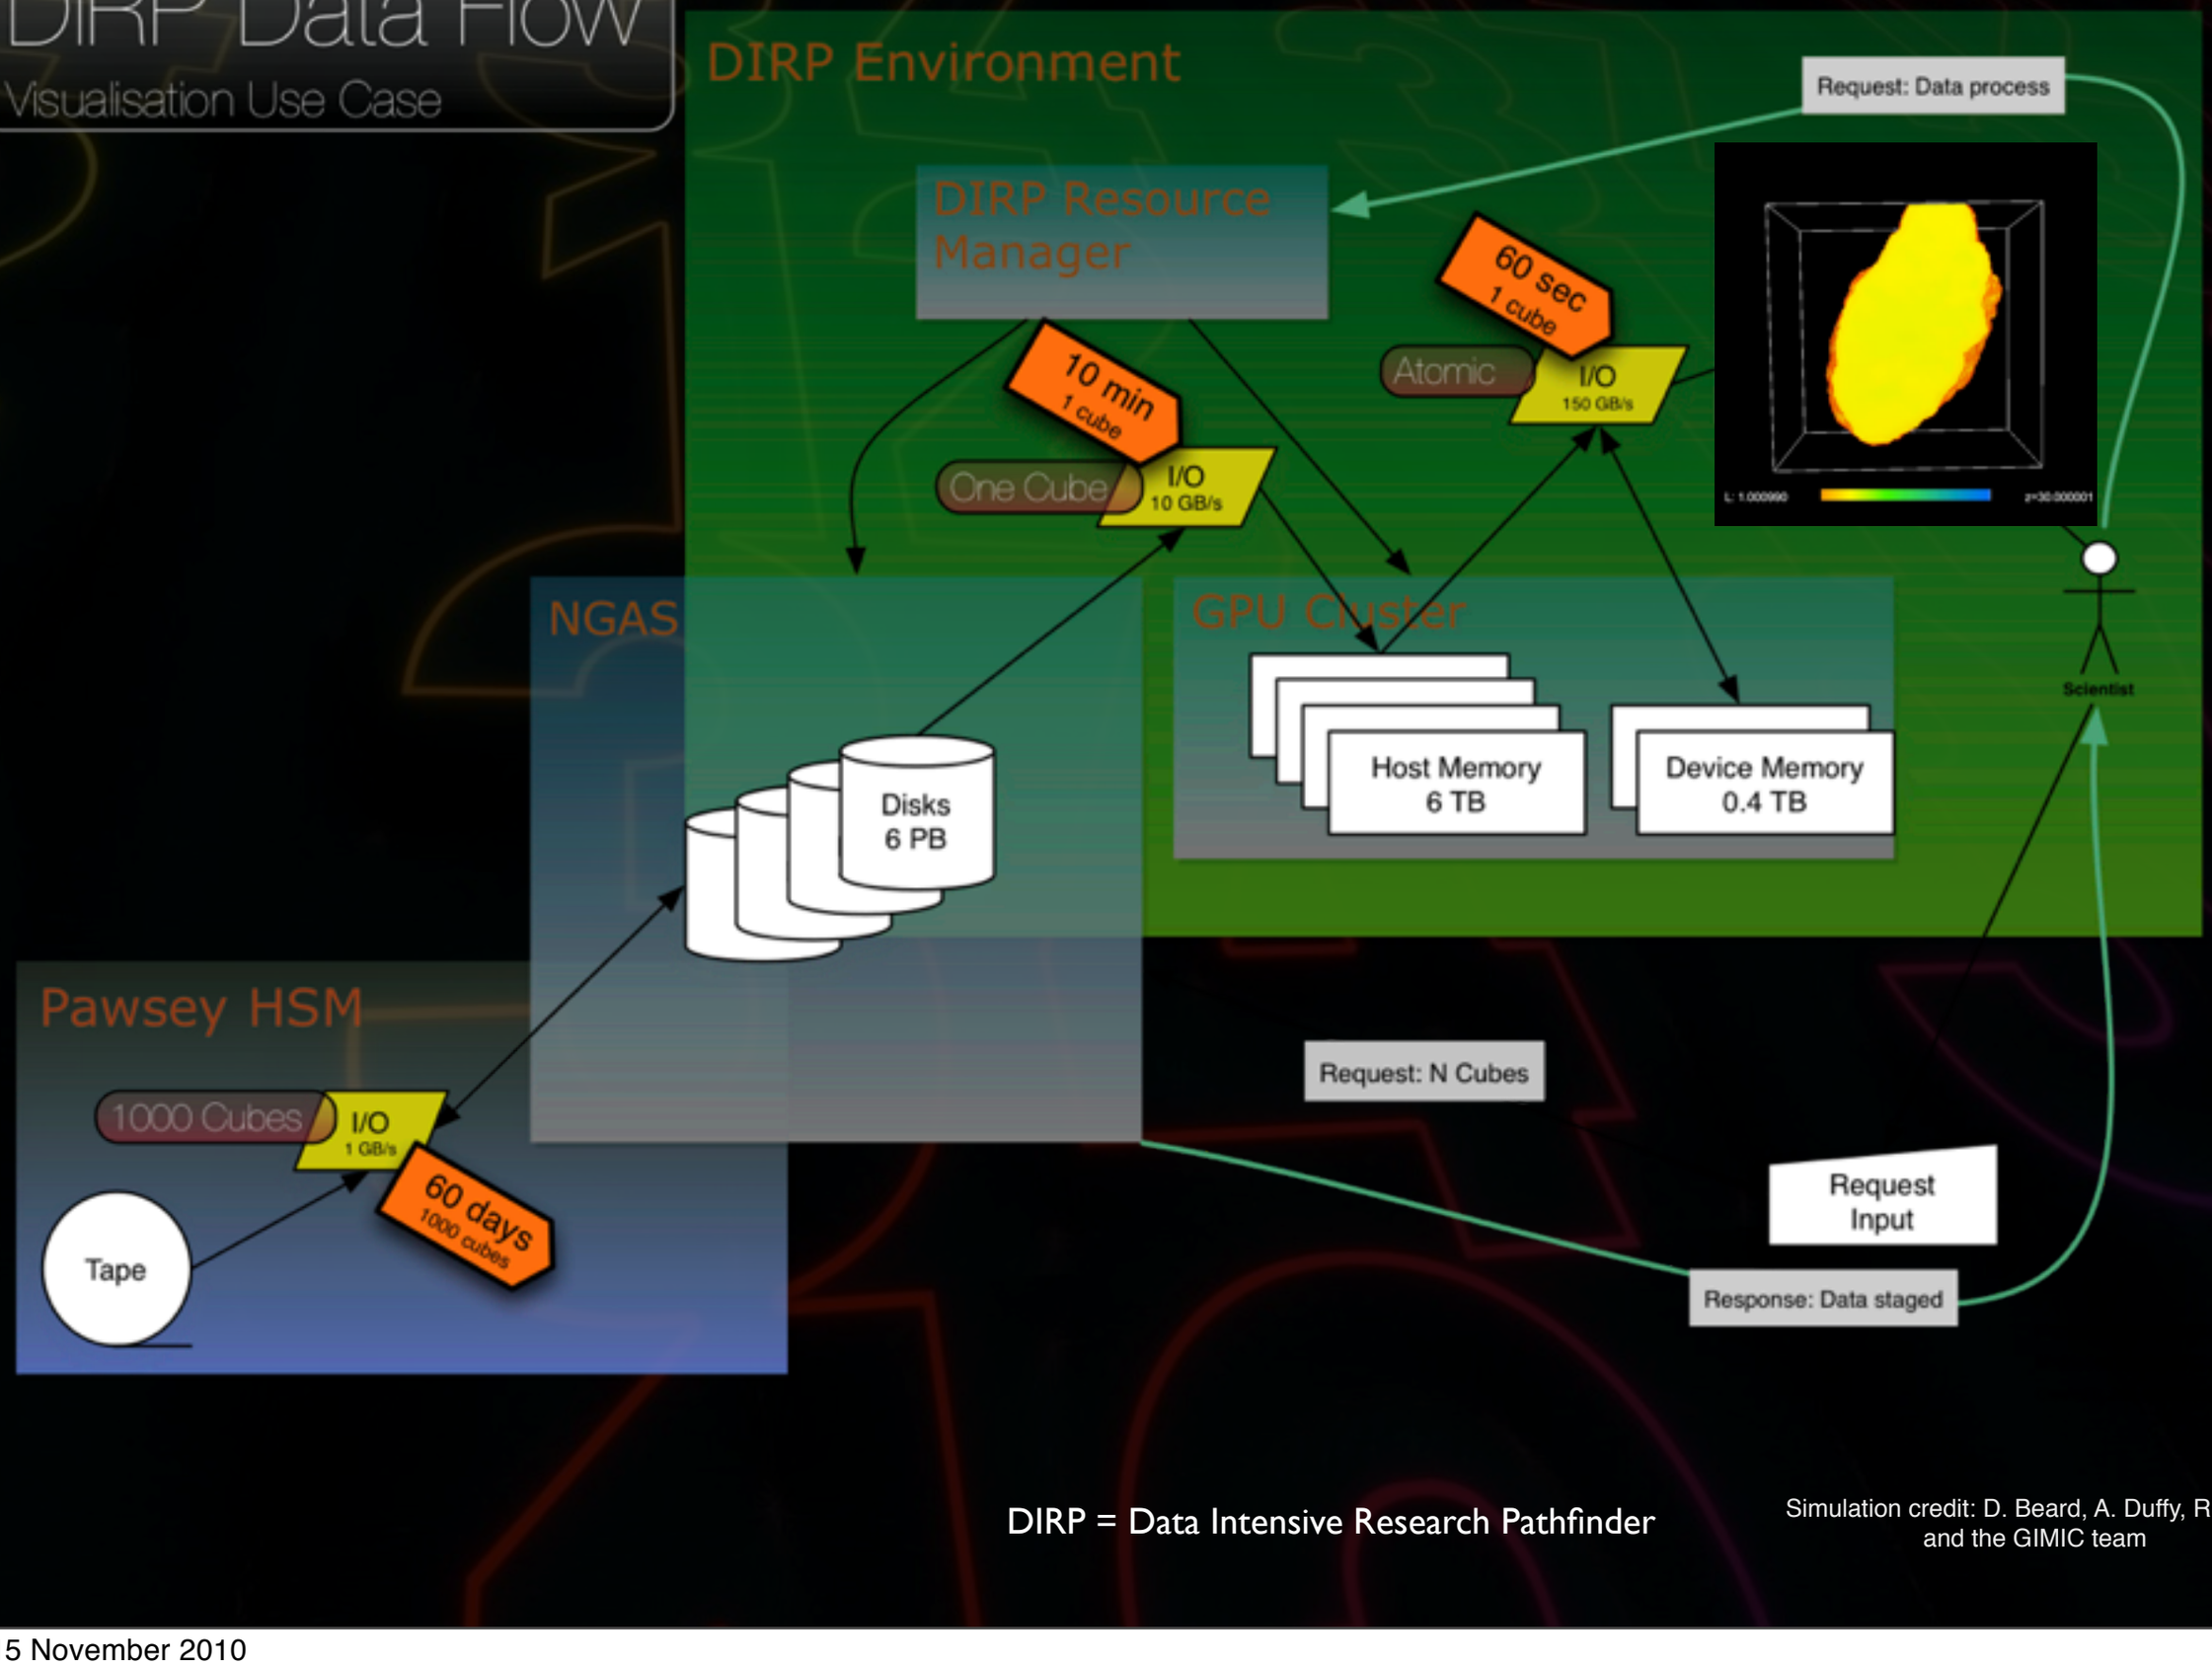

DIRP = Data Intensive Research Pathfinder

Simulation credit: D. Beard, A. Duffy, R. Crain and the GIMIC team

DIR Machine DesignV0.5

- 100 nodes
- 20 TB+ direct attached storage/ node
- 9 TB host memory, 0.6 TB device memory
- 200 GPUs, 1200 CPU cores
- 100 PCI I/O cards or some PCI I/O SANs
- Infiniband interconnect
- Very similar to Johns Hopkins' Data-Scope

DIR Machine Benefits

- 200-500 GB/s aggregate I/O bandwidth between disks and GPUs
- ~ 1,000,000 IOPS on I/O cards/SAN
- > 200 TFLOP/s
- 40 Gbps interconnect
- > 2PB direct attached storage
- Scales very well to bigger installations

DIR Environment

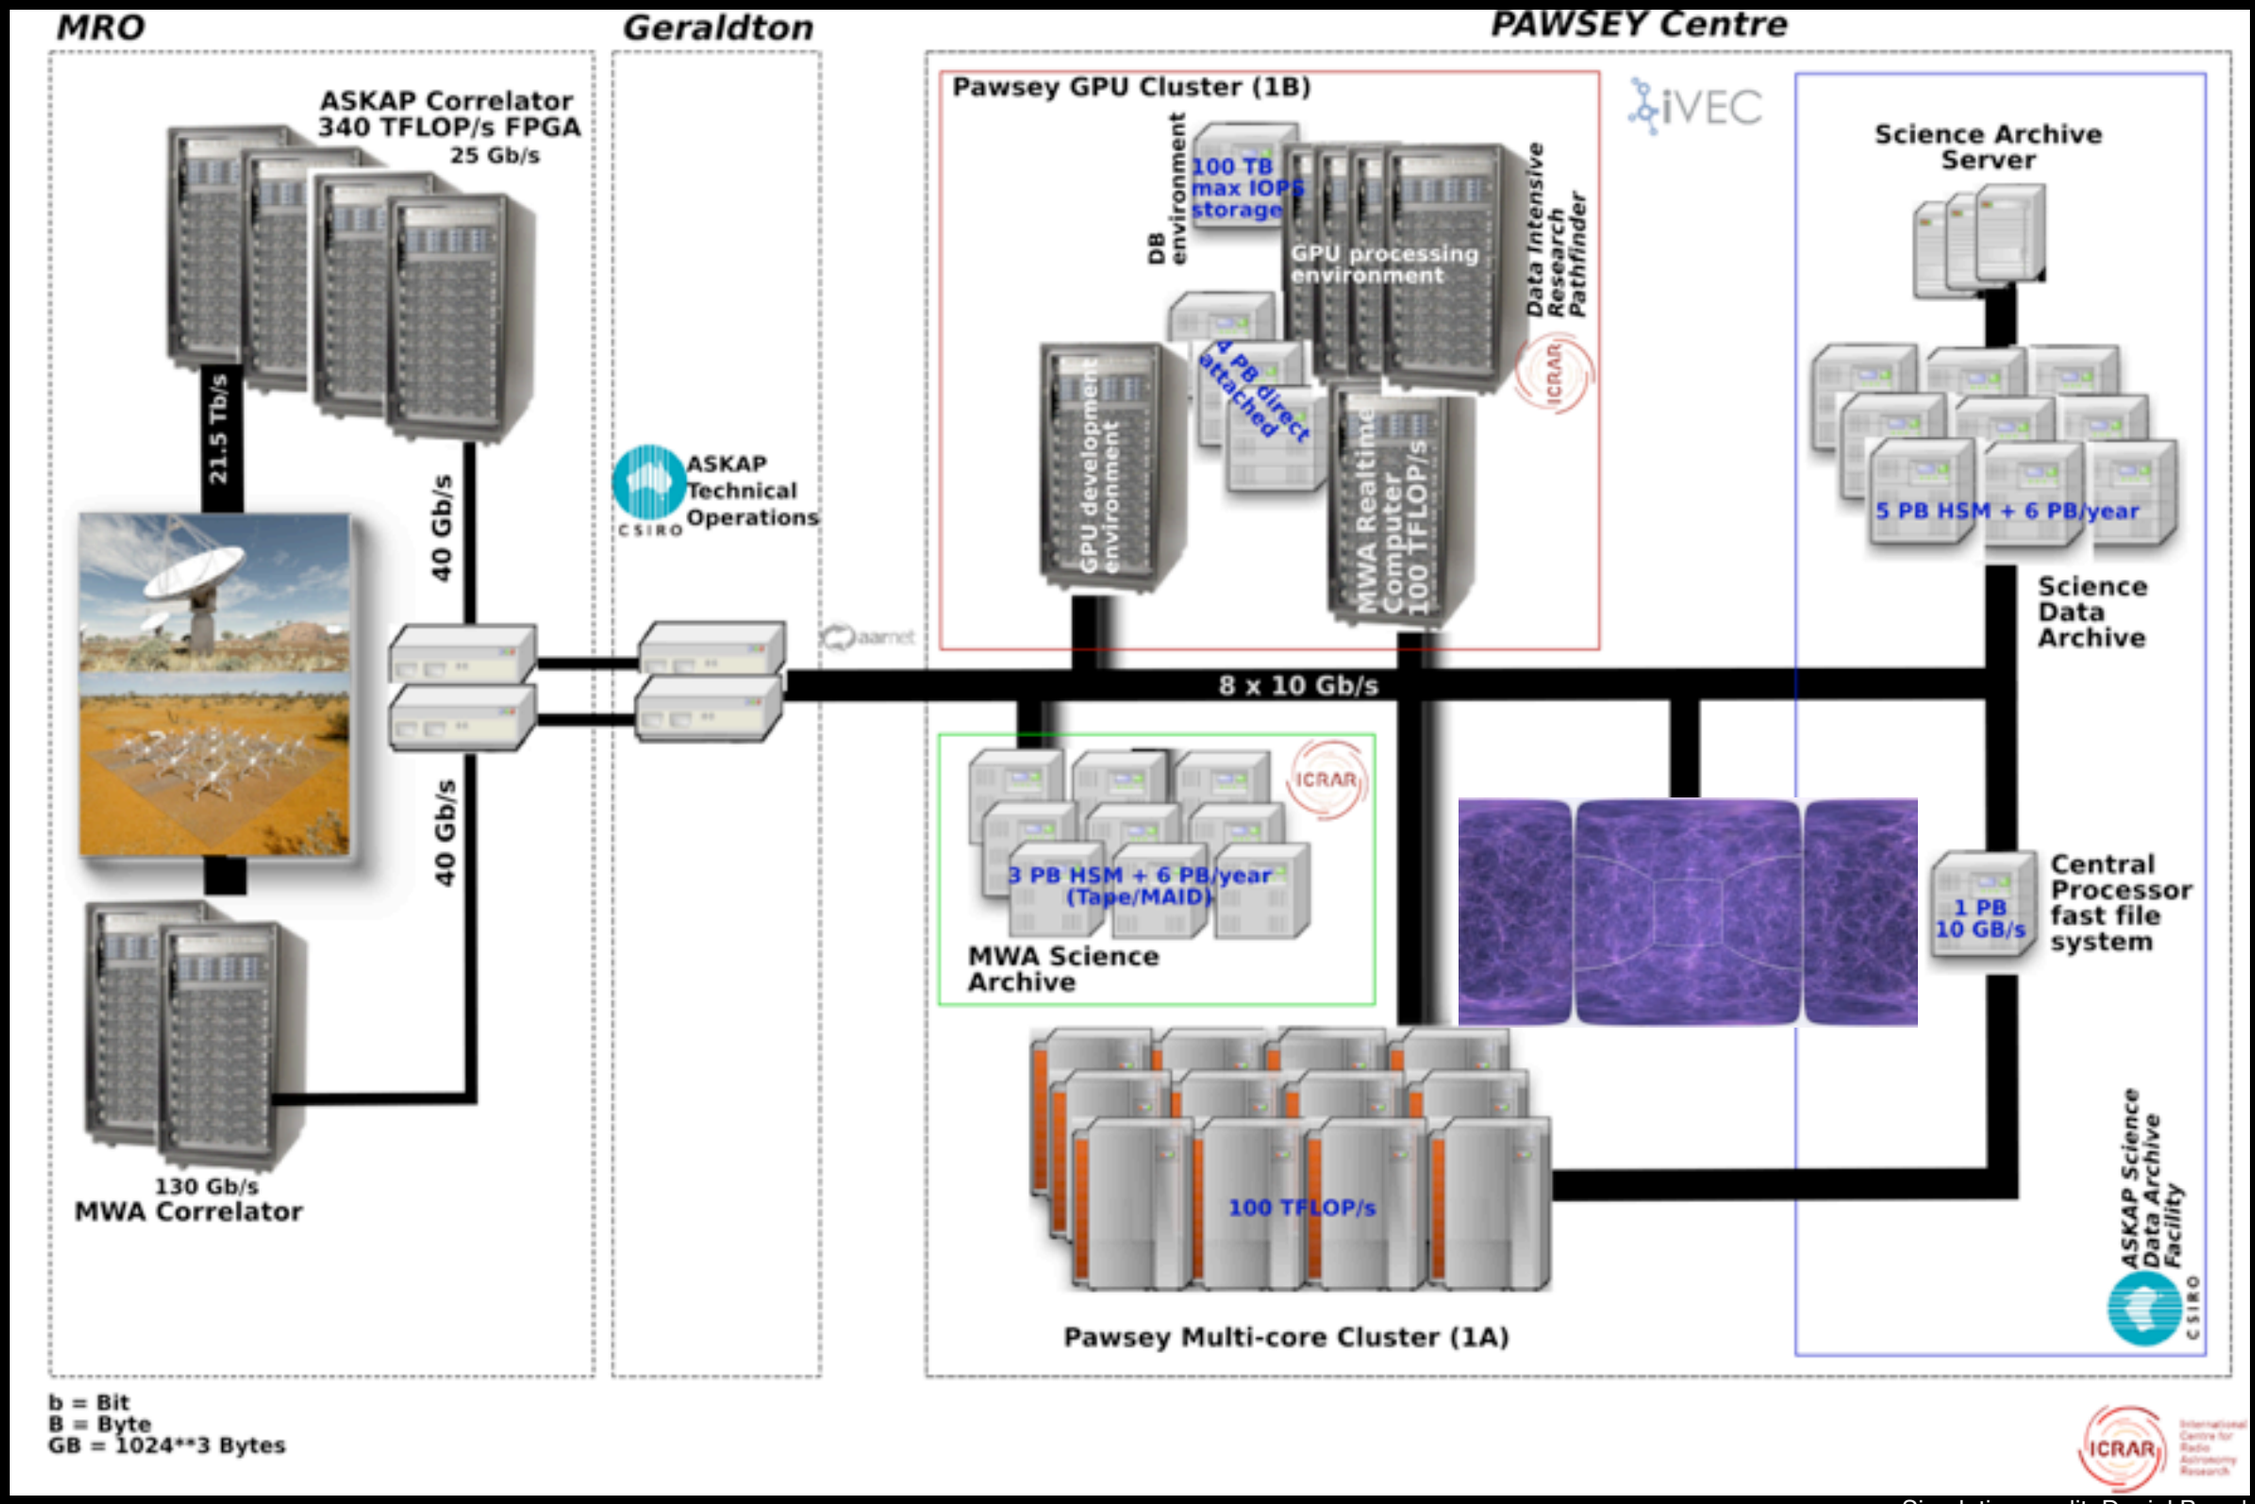

Simulation credit: Daniel Beard, Alan Duffy, Paul Bourke and the OWLS team

International Centre for Radio Astronomy Research

Algorithmic Challenge

Algorithmic Challenge

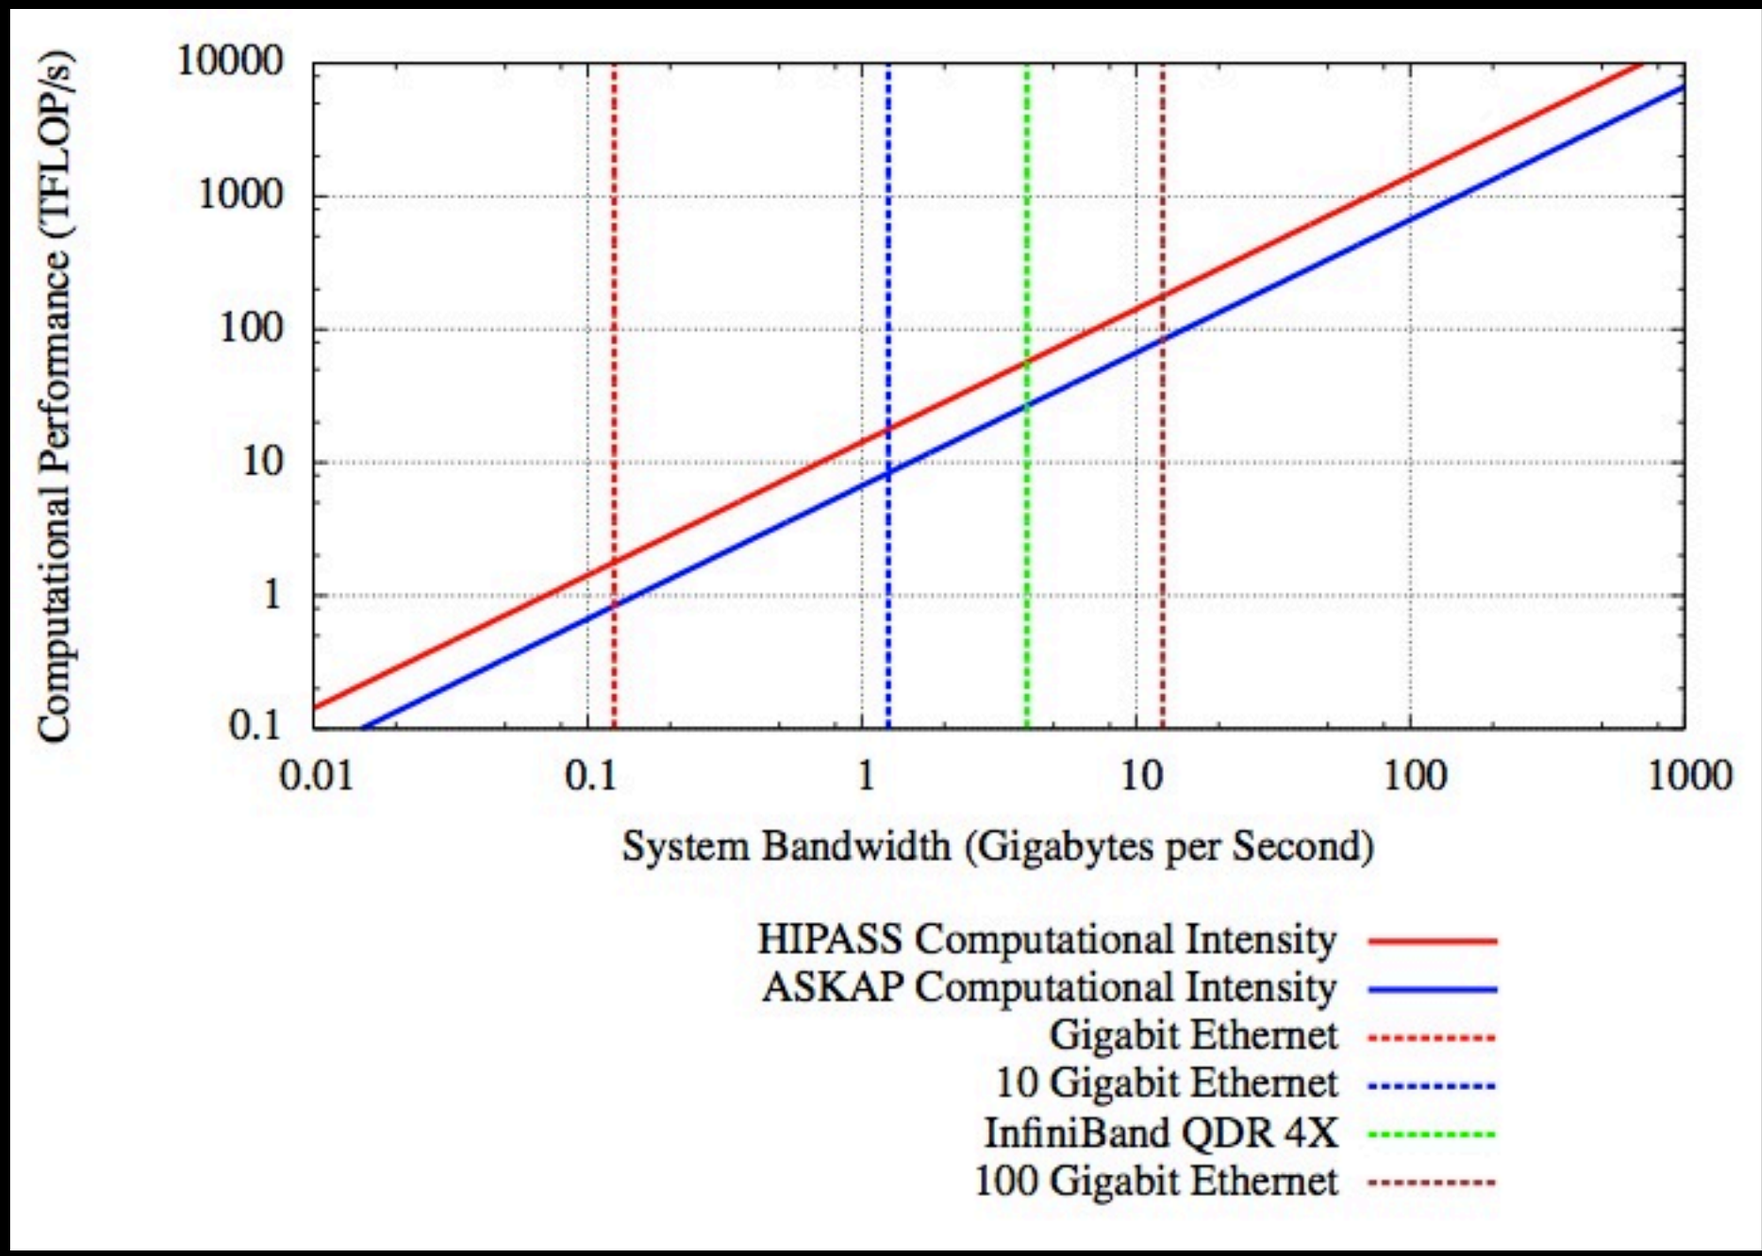

Algorithmic Challenge

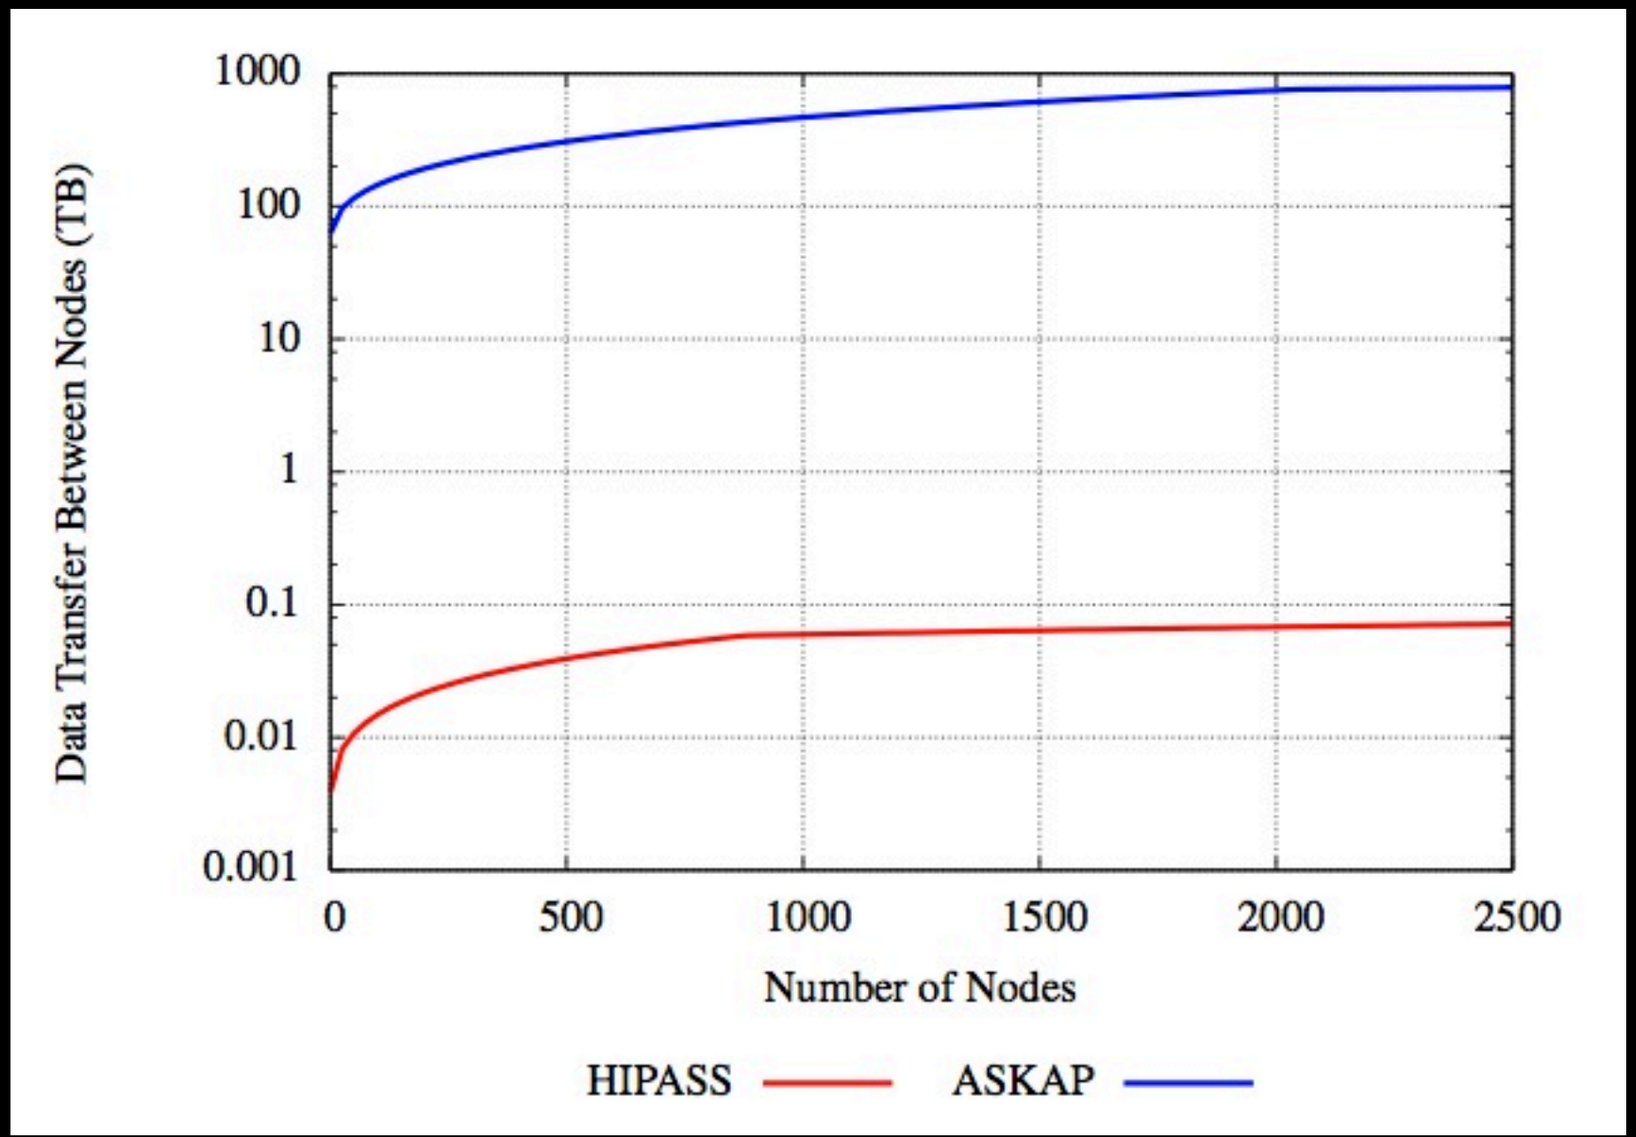

Algorithmic Challenge

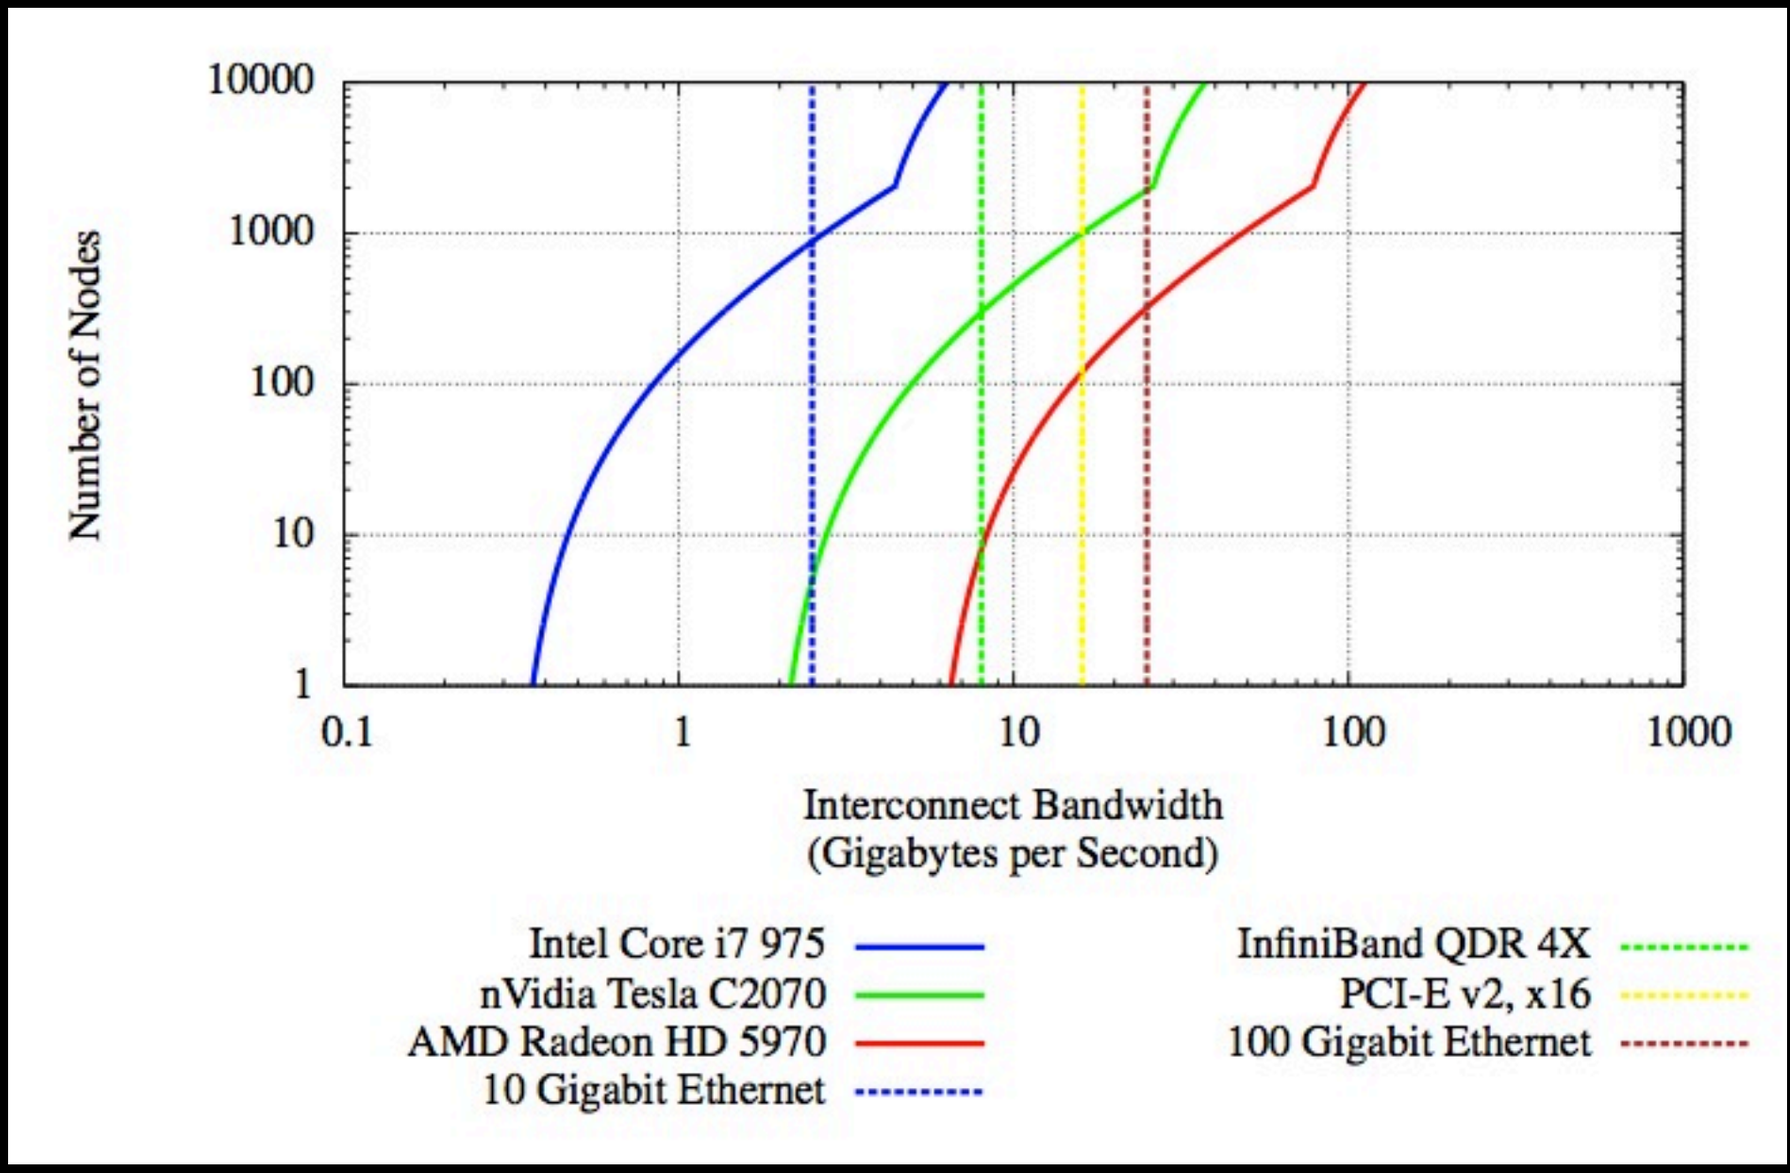

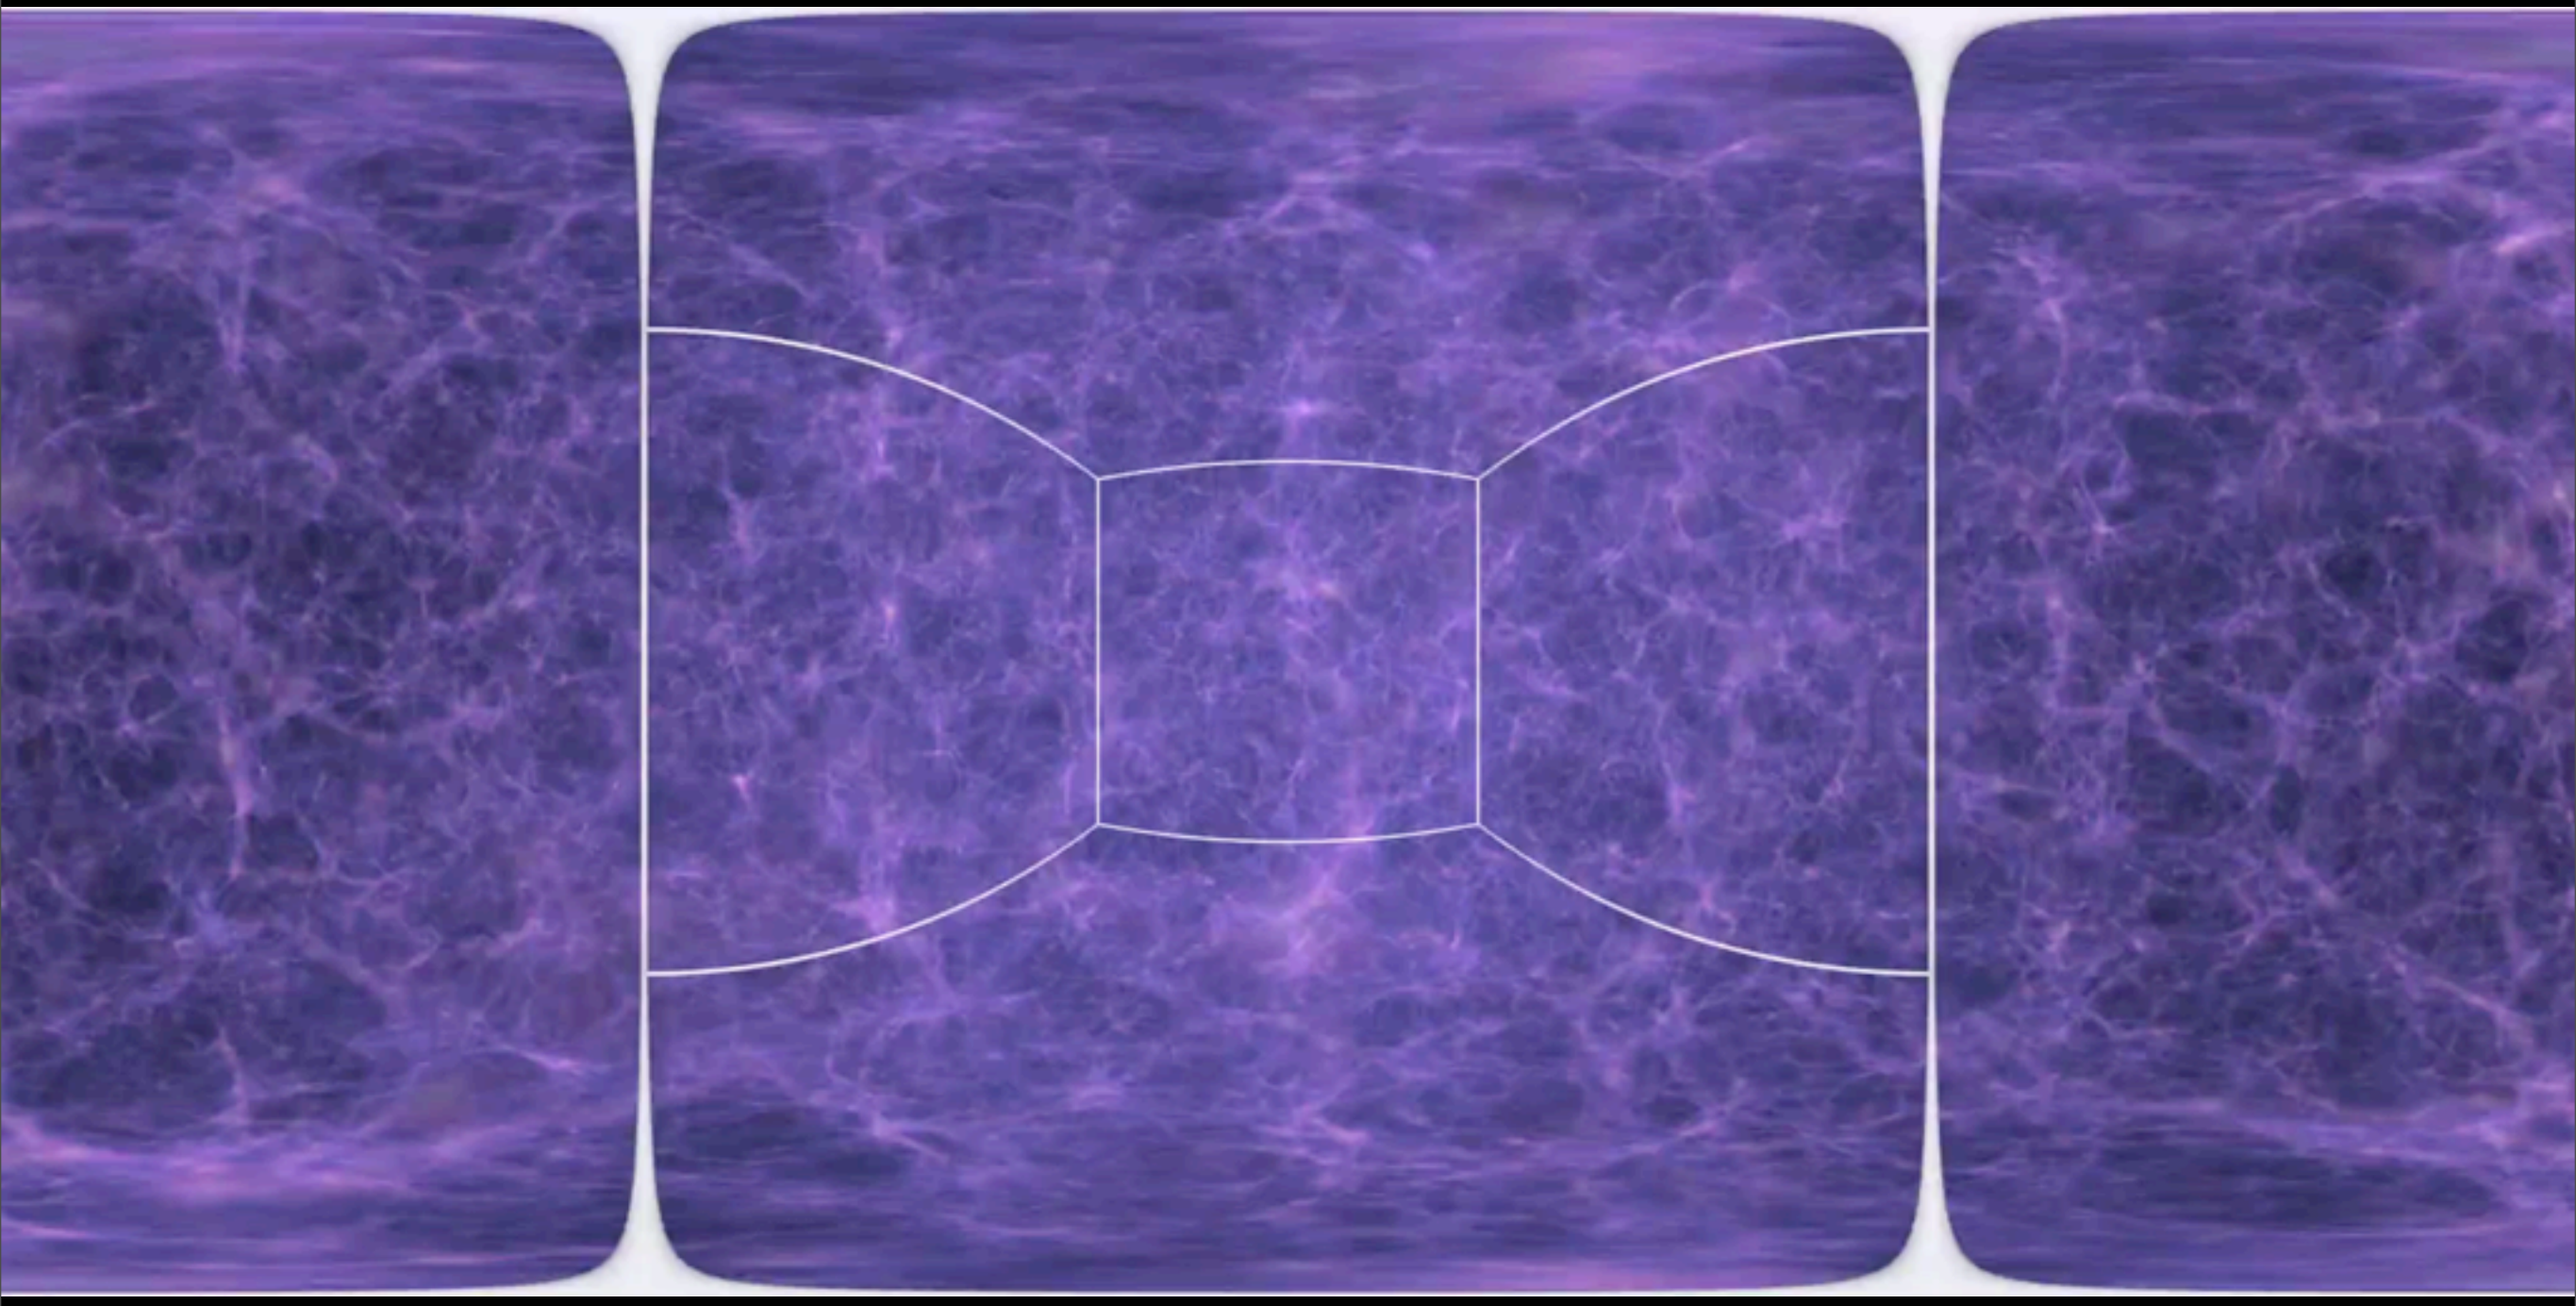

Simulation credit: Daniel Beard, Alan Duffy,
Paul Bourke and the OWLS team

Thank you!

and

Yes, we ARE Hiring!

http://www.icrar.org/employment#hpc

Thanks to Kwan-Liu for the invitation

Thanks to A. Duffy, S. Westerlund, P. Quinn, K. Vinsen, C. Harris and D. Gerstmann for their input.