Data transfer from ALMA to North America

David Halstead, Adele Plunkett

National Radio Astronomy Observatory

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

www.almaobservatory.org

NRAO and ALMA

ALMA is a multinational project with many partners, and three ALMA Regional Centers (ARCs):

- NA: NRAO, Charlottesville, VA, USA
- EU: ESO, Garching (Munich), Germany
- EA: NAOJ, Mitaka (Tokyo), Japan

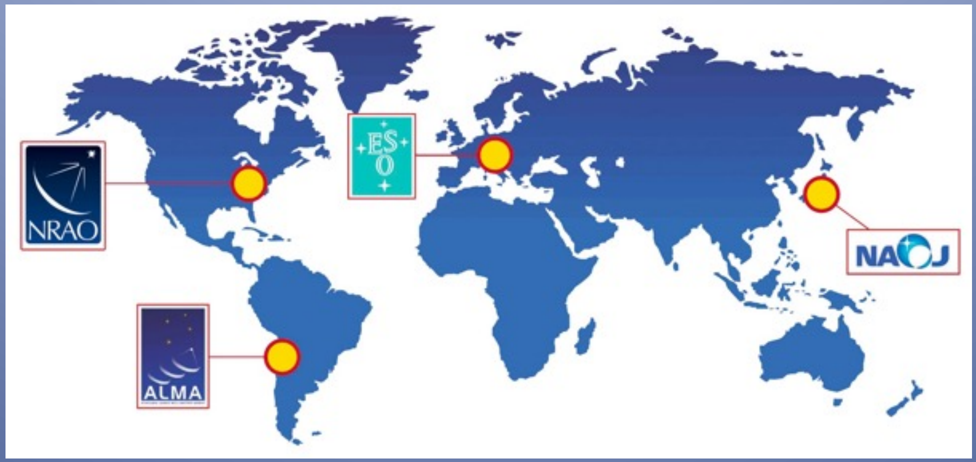

Image: https://science.nrao.edu/facilities/alma/images/arcs.jpg

ALMA telescope

- Largest mm/submm telescope ever built. As an interferometer, it combines signals from multiple antennas to form an image.
- All 66 antennas operational at high site (except for maintenance)
- Sub-arrays possible, and generally 3 projects observed at once:
  - 50 x 12m-antennas (main array)
  - 12 x 7m-antennas
  - 4 x 12m-antennas, observing in “Single Dish” mode (“Total Power”)
- Operated “space mission” style, with pipeline data processing and a science archive at JAO in Chile and each ARC
- First PI projects released to public from the ARCs January 2013 (10-years!)
- *Observing interrupted for ~7-weeks by Ransomware attack Oct 29, 2022*
  - *No data lost, but substantial investment in CyberSecurity during re-build*
- Call for Cycle 10 proposals completed April 2023
- Cycle 10 observing starts October 2023

Radio Interferometry relies on pairs of antennae to emulate a much larger dish

Artificially create a large "dish" using many smaller ones...

... this is called "Aperture Synthesis"

But ... this large "dish" is not a real reflecting surface.....

So how do you make it behave like one ?

... imitate the Physics of a lens.

Radio Interferometry: Relies on pairs of antennae to emulate a much larger dish

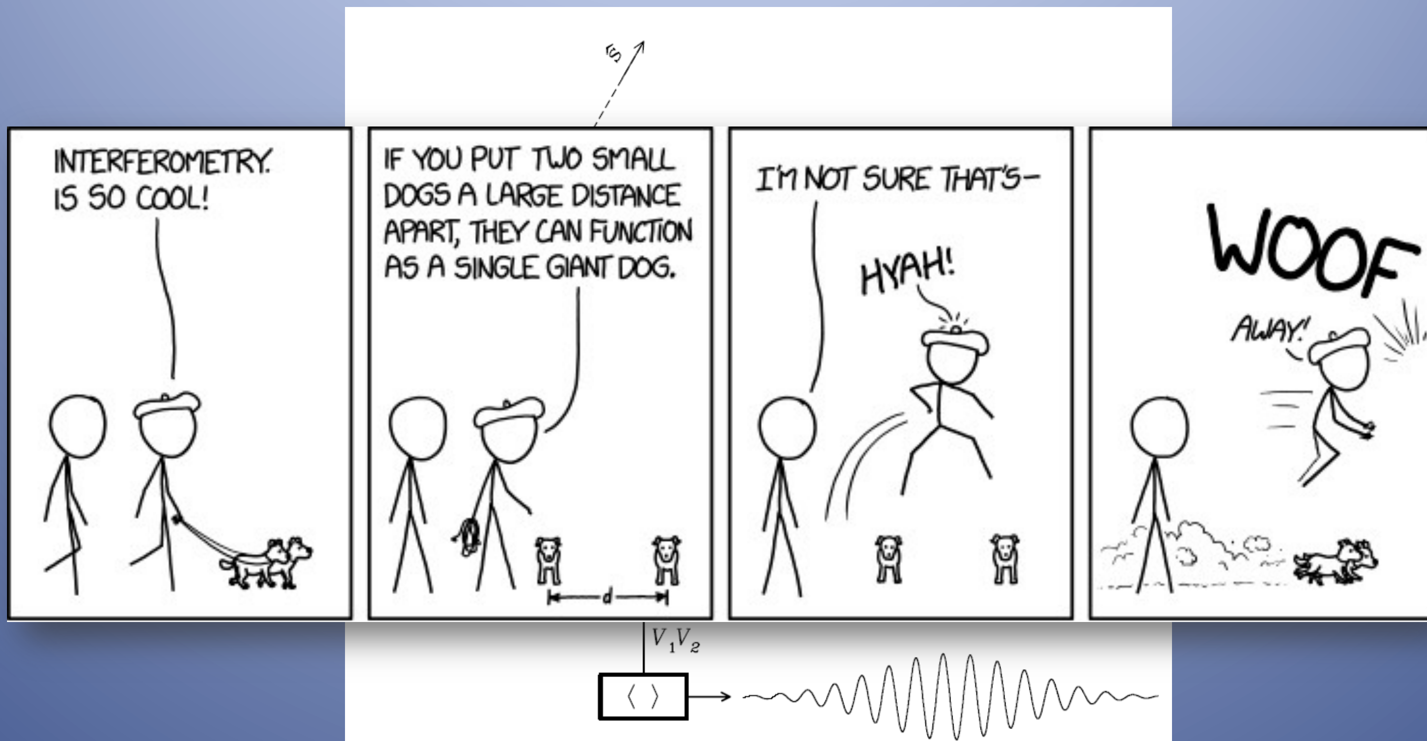

Measure interference fringes

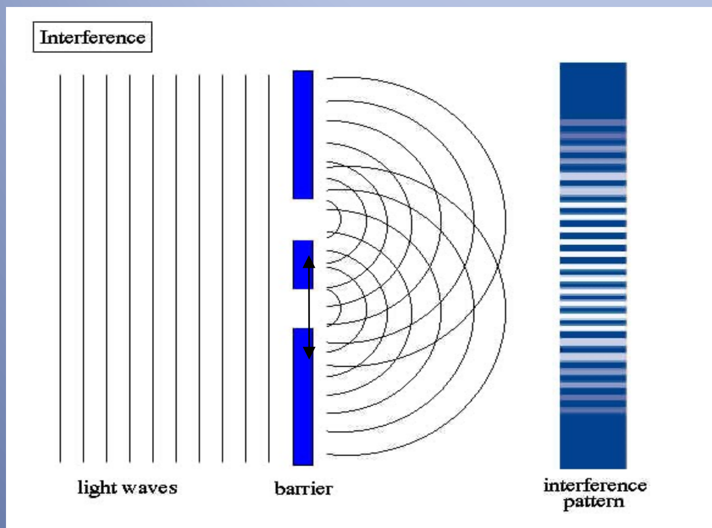

Young's Double-Slit Experiment

Distance between slits controls the wavelength of interference fringes

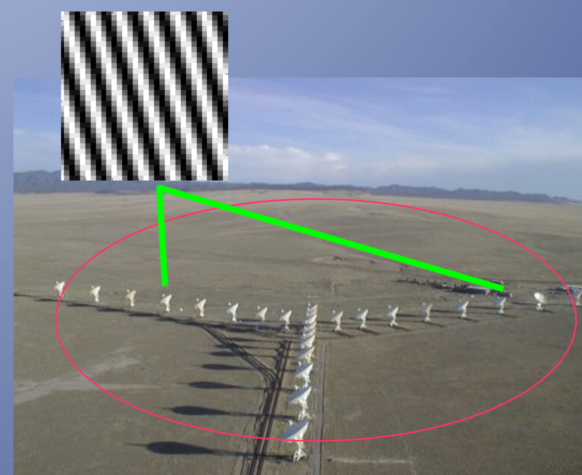

One dish = One slit

Each pair of antennas captures a different 2D fringe

Image Formation

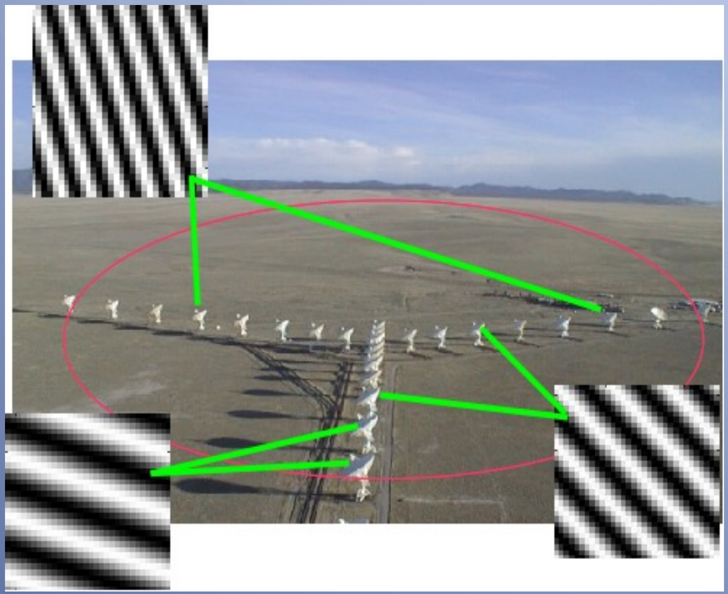

Build an image by combining all measured fringes in a custom built High-Throughput SuperComputer called a Correlator.

2D Fourier transform :
Image = sum of cosine 'fringes'.

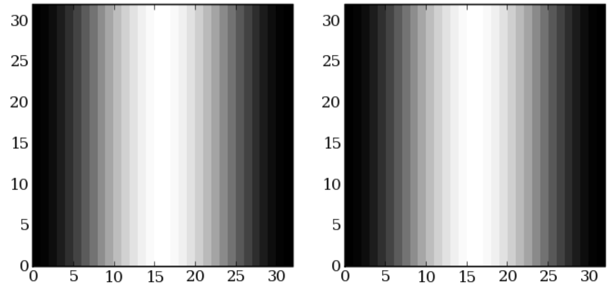

The number of fringe scales with the (number of dishes)^2

ALMA Correlator: HPC@ 16,200 feet

Correlator Quadrant

Tunable Filter Bank Card

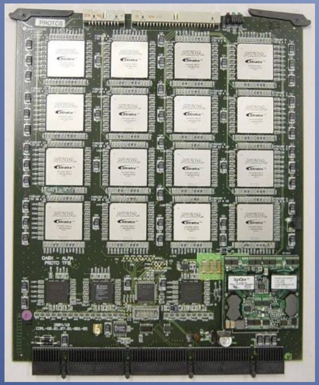

- Receives signals from 50x12m antennae
- 2551 printed circuit boards total in system
- 8192 Altera Stratix II FPGAs on TFB cards
- 32768 custom correlator chips with 4096 processors for multiply-and-add calculations
- Cross-correlation rate 17 Peta ops/sec
- Output specified at 6-60MBytes/sec

Correlator Card

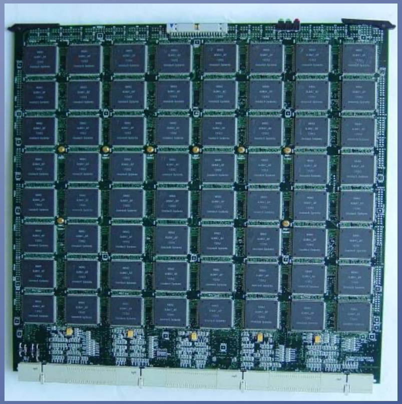

ALMA Telescope Data Flows

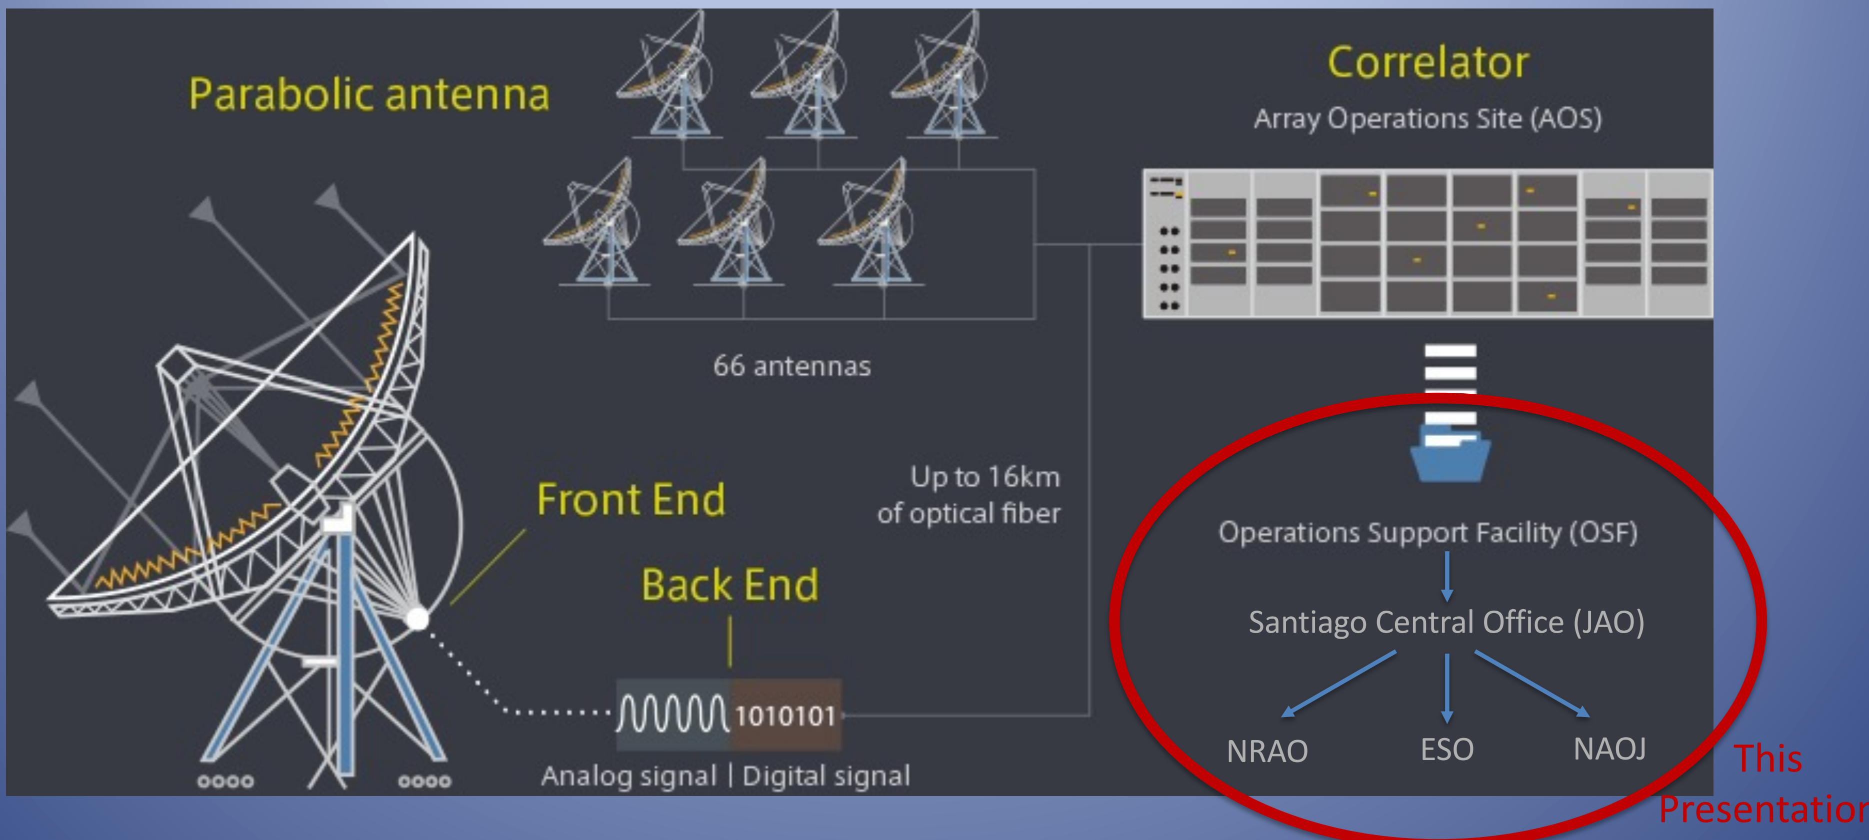

System diagram of ALMA. Various instruments including antennas, receivers, correlators, and data archive work in unison.
Credit: ALMA (ESO/NAOJ/NRAO)

Data transport
Chile to Charlottesville (NAASC)

- MOU between AUI/REUNA for Metro link to SCO
- Santiago to ARCs: individual ARC contracts with REUNA and NRENs.
  - NA: Joint AURA-AUI agreement for 1Gbps committed (burstable to 10 Gbps) from Santiago to I2 via AmLight
  - Link to NRAO HQ from Internet2 through UVa is 10Gb/s
- Typical rate obtained during peak data transfer periods is 2-300Mb/s, with bursts up to 600Mb/s
- Recent transition to most data processing being done in Santiago. Some manual processing at ARCs
- The North American ALMA Science Center (NAASC) hosts the ALMA Archive, and computing for NA users

ALMA Archive Growth (Tera Bytes)

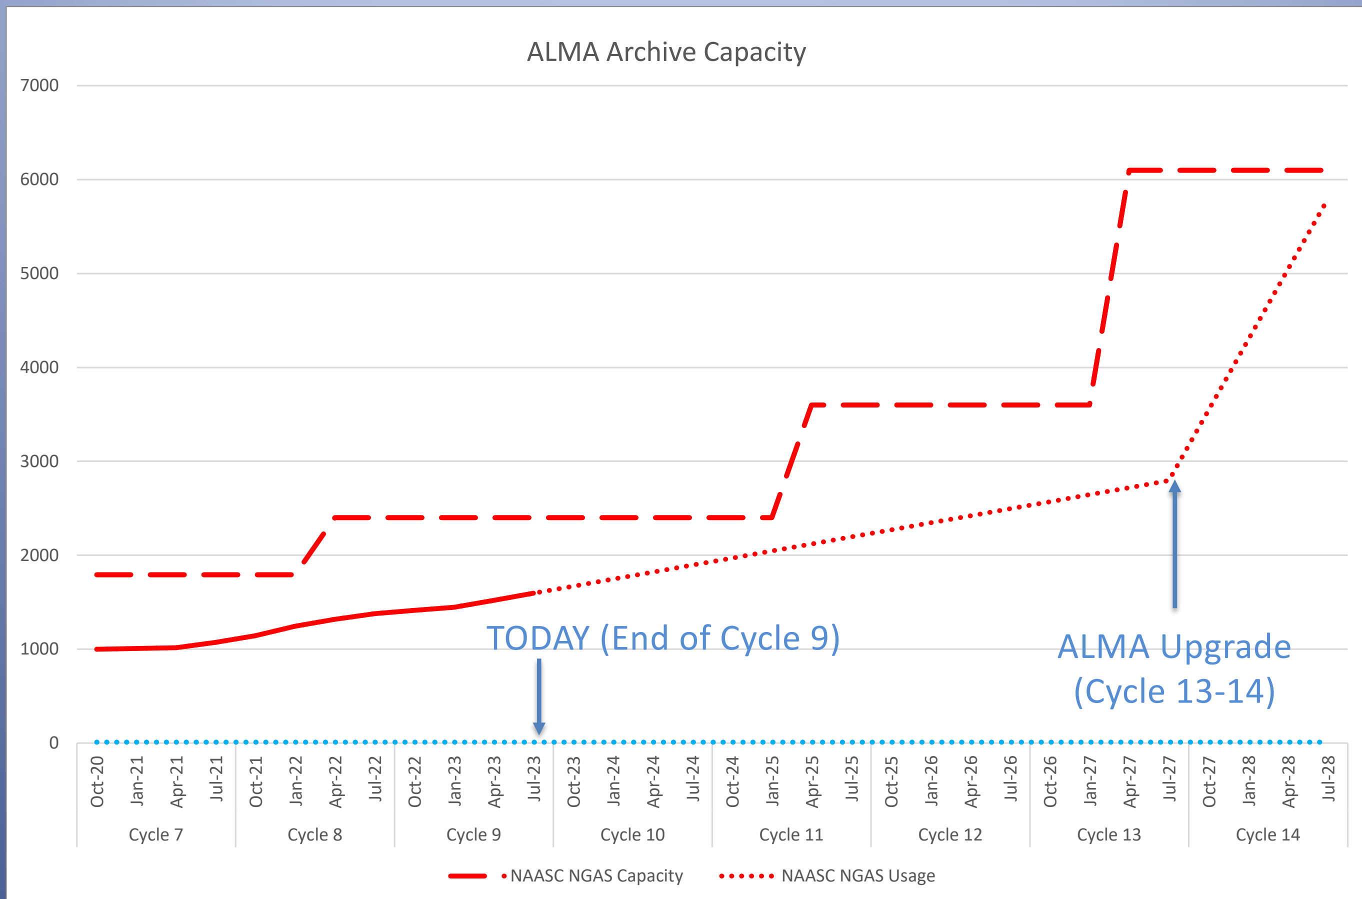

Opportunities

ALMA was a key contributor to the Event Horizon Telescope (EHT) imaging the Black Hole at the center of Messier 87 (M87), 55 million light-years from Earth.

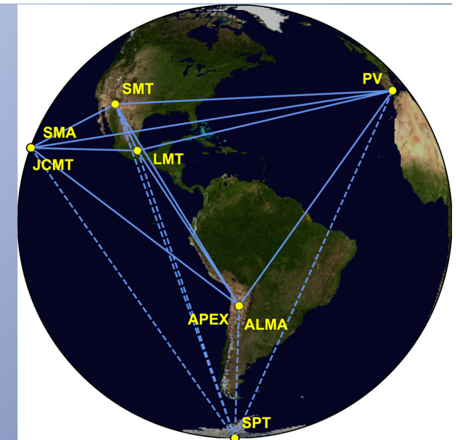

- All ALMA dishes were "phased" to provide a single stream of data for the multi-day, multi-telescope observation BUT…
- Due to bandwidth constraints from geographically distributed telescopes, data was shipped on disk for correlation
- ALSO….
- The much harder imaging of the Black hole at the center of our Milky Way Galaxy, Sagittarius A*

Katie Bouman, S/W developer for EHT, currently CalTech Professor of Computing and Mathematical Sciences (CMS)

2nd Generation ALMA Correlator/Beamformer

- "Advanced Technology ALMA Correlator (ATAC)"
  - Development funding has been allocated (Nov 2022)
  - In line with ALMA Development Roadmap [1], and ALMA2030 vision
  - Likely completion ~2028, as part of ALMA2030 upgrade
  - Specifications available in draft by working group [2]; subsystem working groups currently in progress.
- Deployable location at OSF, physically separate from current correlator and antennas
- Planning for deployment simultaneous with ongoing operations, and with increasing receiver band capabilities over ~15 years.

[1] https://www.almaobservatory.org/en/publications/the-alma-development-roadmap
[2] https://science.nrao.edu/facilities/alma/science_sustainability/Specifications2ndGenCorrelatorV2.pdf

Expected Data Transfer

The assumed average data rate will be (in the 2030s) of the order of 3.6 PB/year - 6.6 PB/year.

- Data transfer from SCO to the ARCs is not identified as an issue when considering the increase in the WSU data volume.
- Most data processing and data transfer to users will be very reasonable; there is a long "tail" to consider of large datasets that fully exploit the new capabilities.
- In 2022, the size of an average downloaded dataset was 20 GB (maximum 2.5 TB). Expect 1TB to 125 TB during upgrade.
  - ALMA project is considering user experience and compute resources (Science Platform) for enabling scientific output.

Disaster Recovery Concern

- The current NA ALMA archive is ~1.3 PetaBytes
- In the event of data loss, the Regional Centers would recover data from the primary ALMA Archive in Santiago
- Over the current 1Gbps link, this would take ~5-6months!
- NA ALMA archive access would be re-directed to ESO during any outage
- Additionally, recovery of the Oracle DataBase would take many weeks

Summary

- Ramp-up of the ALMA data rate has been slower than anticipated, allowing us to stay ahead of the curve.
- Data flow mostly from JAO to ALMA Regional Centers, with data processing mostly at JAO.
- Still monitoring how the network performs when transferring ~10TB/day in multiple parallel streams.
- We must establish a link with 10Gb/s of dedicated bandwidth within the next 1-2 years to improve transfer speed to and from Chile for bulk reprocessing, and to help with occasional large data and metadata transports (e.g. a DB export).
- Most new developments (e.g. next generation correlator) on ~5yr timescale can probably be accommodated without increasing the data rate by more than a factor ~4.

Questions?

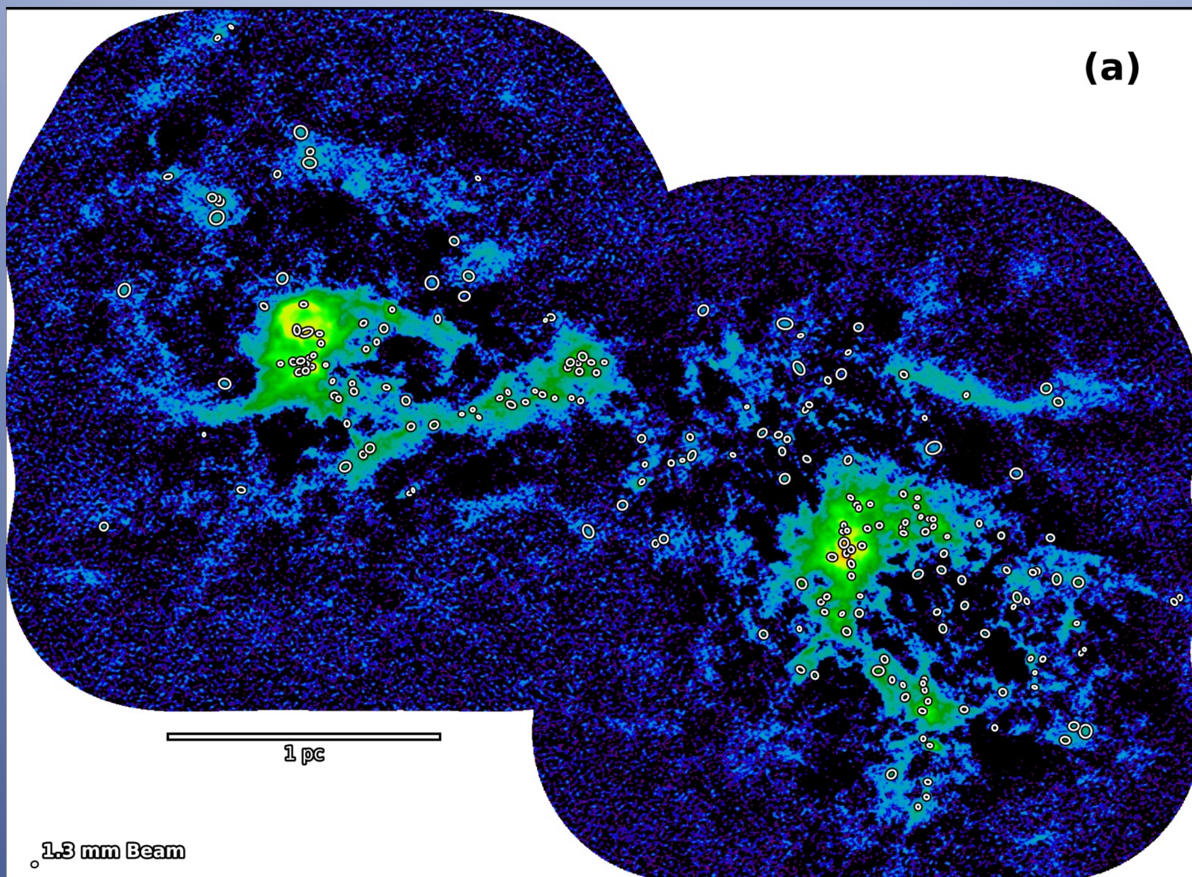

Mini ”star-burst” studied by Pouteau et al. (2022), based on ALMA Large Program data