# Taming the ALMA data avalanche

Felix Stoehr
Subsystem Scientist of the ALMA Science Archive


IAU General Assembly, SpS15

Beijing, 28.8.2012

# Outline

- ALMA
- The data challenges
- Outlook
- Who? Too many to name them on a slide. The credit goes to the whole ALMA team.

# ALMA

# ALMA First Light

# ALMA IN A NUTSHELL

- 66 antennas at 5000 m elevation in the Atacama Desert
- Interferometry with baselines up to 16 km
- Wavelength range from 3 mm to 400 μm (84 to 702 GHz)
- Built by ESO, NRAO, NAOJ in cooperation with Chile
- Timeline:
  - Early Science Cycle 0 observing is under way
  - Early Science Cycle 1 Call for Proposals very successful
  - Cycle 1 observing around beginning of 2013
  - Completion of the full array: second half of 2013

# The data challenges

1. Data rate: 200TB/yr with peak rates of 10x average
2. Data storage: "All data taken have to be stored forever"
3. Archive copies: full copies on four continents
4. Operation: general-purpose telescope
5. Pipeline: science-grade products for all scientific data
6. Support: computing power and hands-on help for scientists

# 1 Data rate

- The **theoretical maximum** data rate of ALMA correlator is 15PB/yr = 512MB/s = 4096 Mb/s (currently 1/8$^{th}$ but upgrade is easy)
- Based on the predicted data rate requirements of the **expected type of science observations** the official average data rate has been set to 200TB/yr = 6.6MB/s = 50Mb/s
- The NA ARC (M. Lacy, D. Halstead) updated these values with the **experience gained through Cycle 0**. The expected data rates are higher: up to 700TB/yr = 23MB/s = 190Mb/s
  https://science.nrao.edu/facilities/alma/naasc-memo-series/naasc-memos/110.naasc-data-rates
- The **peak data rate** is 10 times the average rate and the system was designed to support this rate for long periods

# 1 Data rate

- **Choice**: the system was designed to contain no hard bottlenecks. Most systems are build out of **commodity hardware**. It can grow naturally to support higher rates if needed.
- **Lessons learned**: so far, this seems to work out very well
- **Way forward**: monitor the actual usage very closely and upgrade as required

# 2 Data storage

- ALMA storage solution requirements:
  - Cost-effective, scalable, future-proof, PB-scale
  - High read/write performance despite large variation in file sizes (a few kB up to hundreds of GB)
  - Large I/O speeds to the pipeline processing
  - Built-in file management layer (database)
  - Automatic consistency checks
  - Support globally distributed archives

# 2 Data storage

- **Choice**: NGAS http://www.eso.org/projects/dfs/dfs-shared/web/ngas
  - Developed by ESO for the VLT
  - Combined hardware/software solution
  - Commodity-hardware approach
  - Provide online processing capabilities
  - Portable
  - Powerful plugin architecture (archiving, processing, ...)

# 2 Data storage

- **Lessons learned**: ALMA was able to **reuse** the existing **NGAS** solution and ESO had substantial experience running it. However, improving the I/O performance for small files, writing the multi-stream mirroring and making it robust to network failures was considerably more work than anticipated.

- **Way forward**: continuously **increase storage capacity** matching the actual data rate. Avoid (external) vendor lock-in to allow to adapt smoothly to new technologies.

# 3 Archive copies

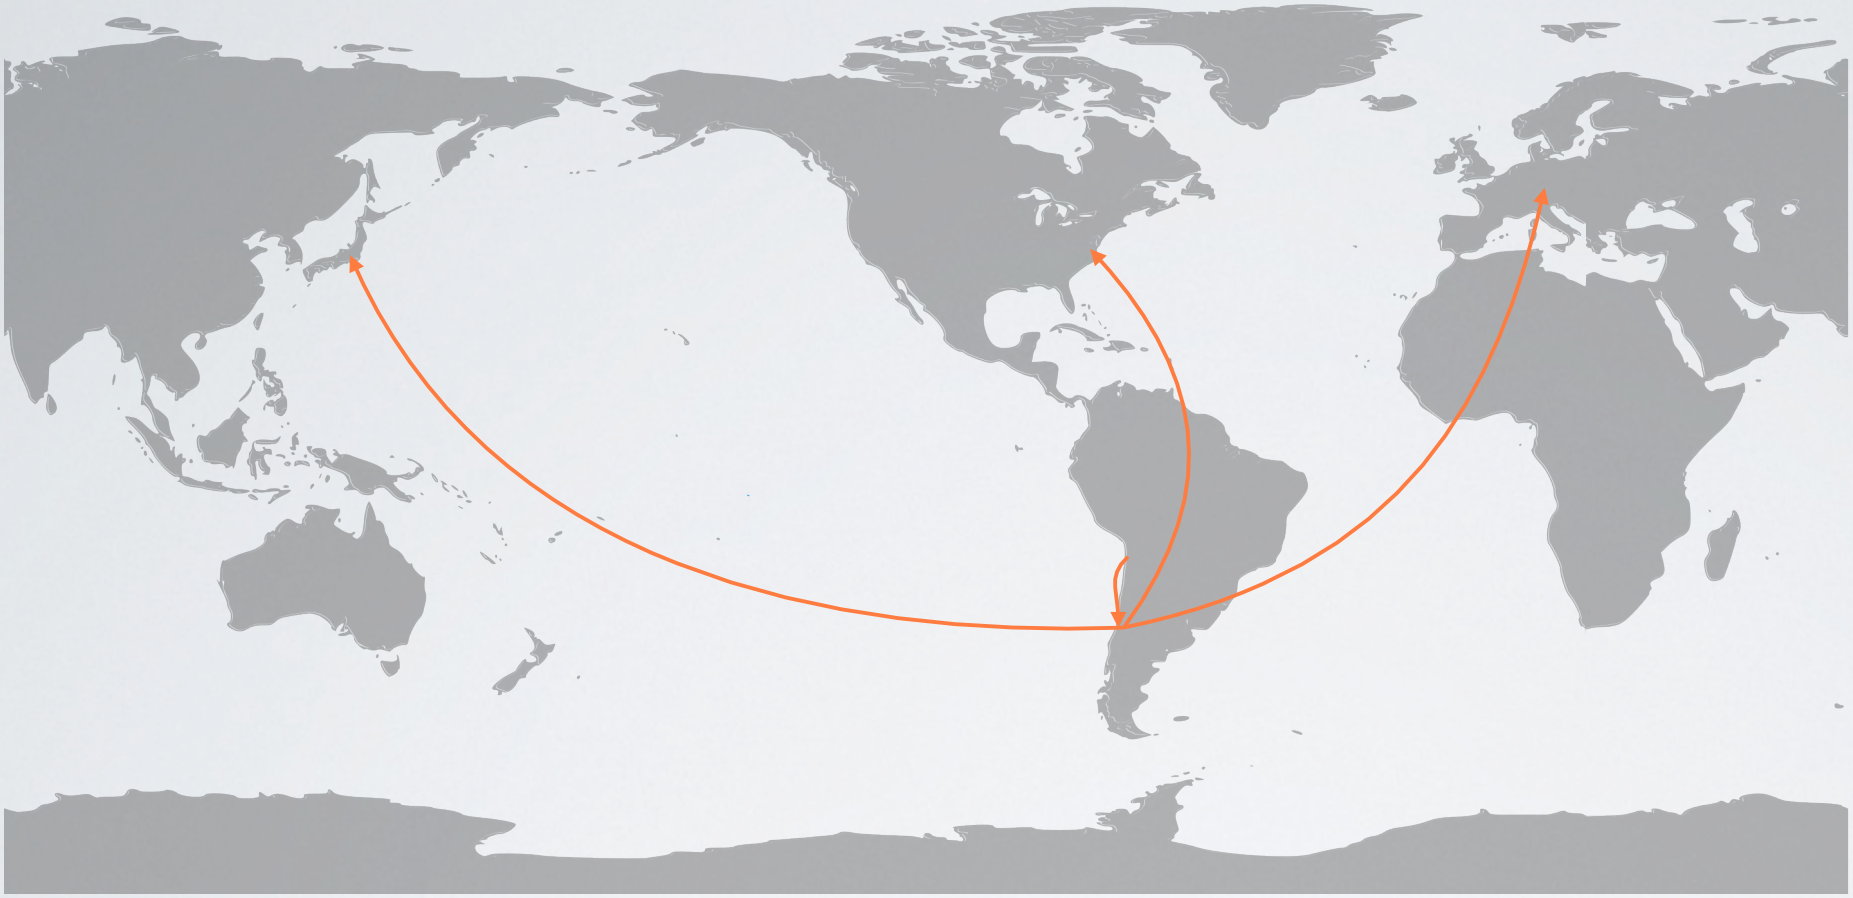

# 3 Archive copies

- The main archive is in Santiago de Chile
- Full copies in Charlottesville (NA), Mitaka (EA), Garching (EU)
- The copies serve as **backup** of the main archive and allow the ARCs (ALMA Regional Centre) to provide user-support, reprocessing, quality control, phase 3
- P2P proposal of the NA ARC (M. Hatz, M. Lacy, D. Halstead) to **reduce network transfer cost**

# 3 Archive copies

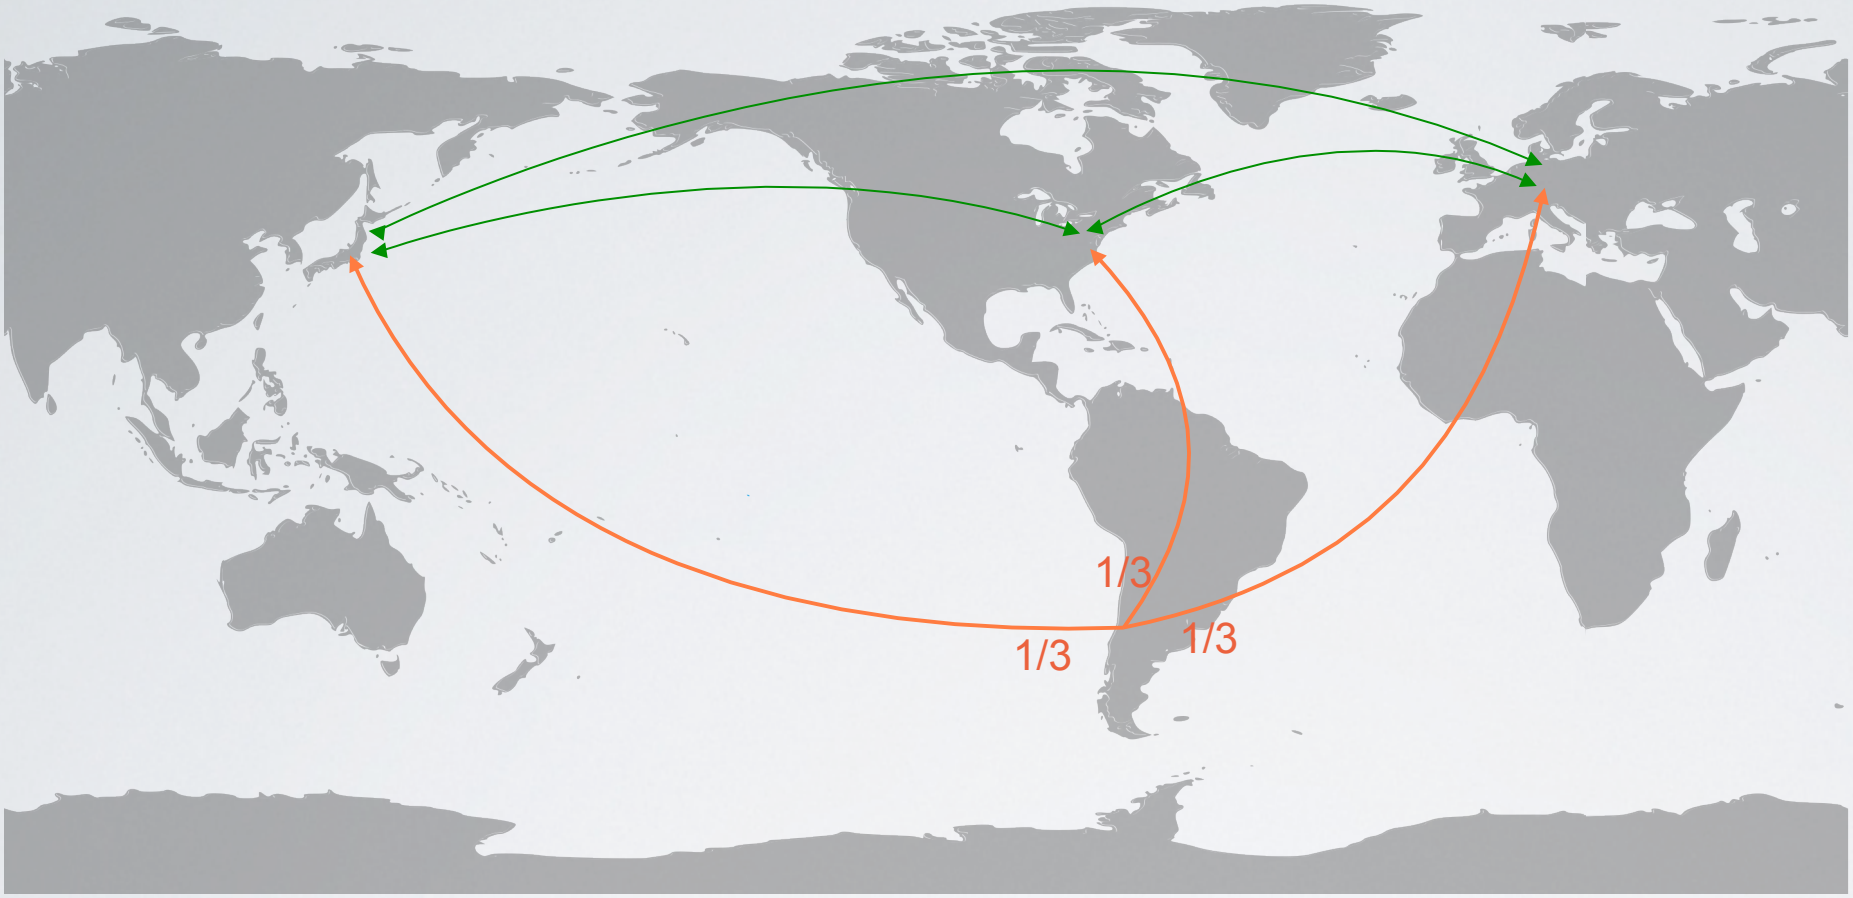

# 3 Archive copies

- **Choice**: VPN **network transfer** instead of disk shipping
- **Lessons learned**: also from past experience at ESO we know that wherever possible, network transfer should be used. Individual files arrive much faster, it is much less work-intensive, much more robust and in the end more cost-effective.
- **Way forward**: implement the NA proposal to **reduce network** cost in Chile

# 4 Operation

- ALMA is a **general purpose telescope** with standard calls for proposals once a year
- ALMA will support a large variety of **different types of observations** (continuum, line, spectral sweep, mosaic, solar) with different observing settings (12m array/ACA/TP) in different configurations

# 4 Operation

- ALMA work-flow (courtesy Erich Schmid)

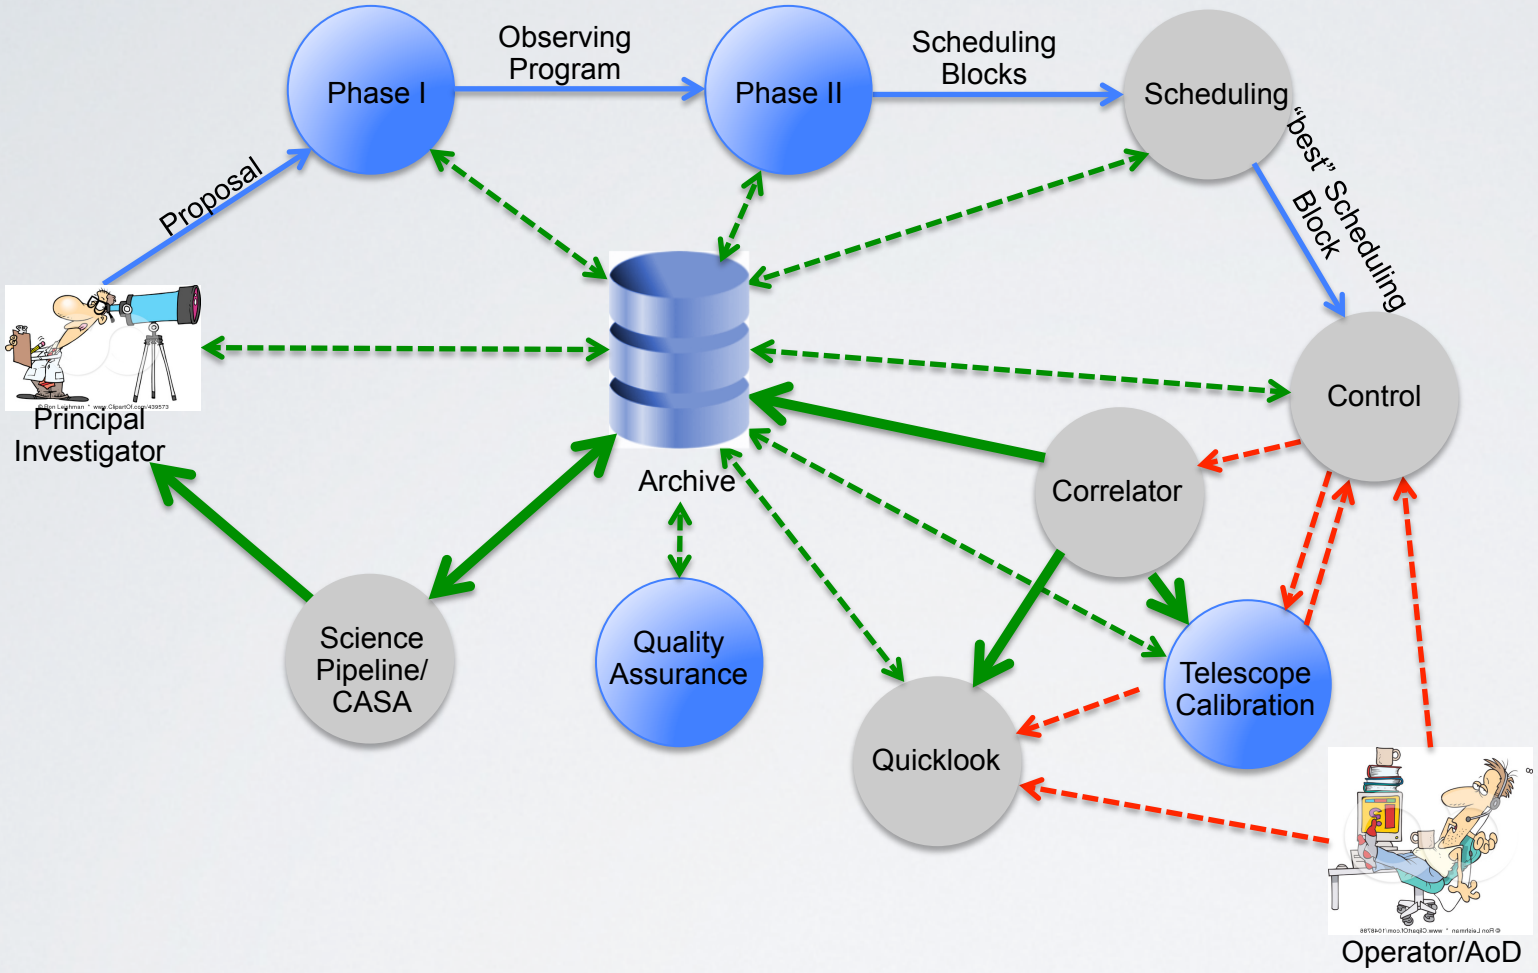

# 4 Operation

- **Choice**: fully integrated system with the Archive in the center using all the same interfaces has been built. A sophisticated **observing tool** and **science data model** (ASDM = ALMA Science Data Model) have been developed.
- **Lessons learned**: so far, we think that the choices were the right ones. The performance of the XML metadata database turns out to be rather low.
- **Way forward**: The ALMA Science Archive including VO services is being built. Metadata will be harvested into denormalized, search-optimized relational tables. In the longer term we will move at least parts of the metadata to a relational database.

# 5 Pipeline

- The amount of data of a single project can easily exceed the data-reduction resources available to a typical user today
- A fully **automatic pipeline** is required for ALMA delivering **science-grade** products (a first in radio astronomy)
- Policy: All ALMA science data will be **reduced by the project**

# 5 Science grade?

- What do PIs do?

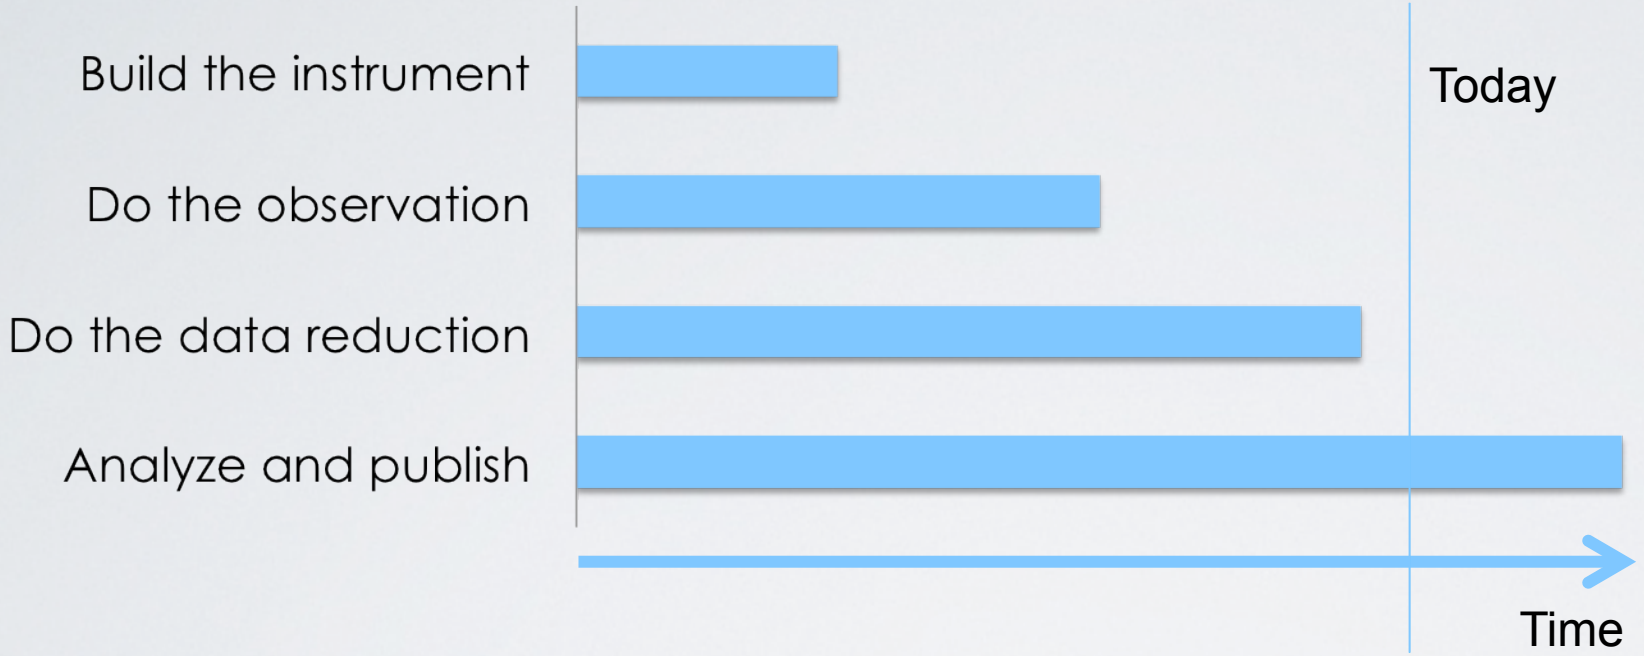


- As telescopes will become more complex and deliver more data, telescopes will move to deliver **science-grade** products

# 5 Pipeline

- **Choice**: An **analysis software** package CASA (Common Astronomy Software Applications) was written (nearly) from scratch. A **pipeline** (heuristics + CASA) was developed.
- **Lessons learned**: so far this approach looks very good
- **Way Forward**: Commission the pipeline and bring it into operations.

# 6 Support

- If users want to re-analyze the data e.g. an optimized reduction for their particular science case, they **need help**.
- ALMA is meant to be a telescope that can be used by **all astronomers**, not only those that have a radio background
- **Choice**: the three ARCs provide **face-to-face support**.
- **Lessons learned**: already now, despite the small data sizes, the face-to-face support is highly appreciated.
- **Way Forward**: provide a visualization solution (development plan). Provide the possibility for users to submit pipeline jobs with different parameter settings on the ARC processing clusters.

# Long-term evolution

- ALMA operations are expected to last **many decades**. Scalability and flexibility in hardware and software are extremely important.
- For the areas that do not scale with Moore-type laws, **parallelization/clustering** is needed.

| | Moore-type law? |
|---|---|
| Hard disks | yes |
| Bandwidth | yes |
| Single TCP stream | no |
| CPU power | yes |
| Single-core speed | no |
| Database rows | no |

# Impact of science on data rates (and vice versa)

- Science influences data rates
  - We do not know which data rates we will get: They **depend hugely on the science** that gets accepted
  - The new analysis shows that scientists are requesting **more wavelength ranges** to be observed simultaneously than anticipated and on-the-fly interferometry which is very data-intensive is not yet implemented

# Impact of science on data rates (and vice versa)

- Data rates influence science
  - Probably many/most scientists will **not reduce the data themselves** any more or at maximum tweak the imaging
  - Extremely large amounts of 3-4 dimensional data available (6mas resolution at 675GHz, up to 7680 channels, full polarization, up to 32 spectral windows simultaneously)
  - The data is **very well described** which (hopefully) will allow for very **powerful searches** (data mining)

# Summary

- High data rates and the long-term nature of ALMA are a real **challenge** and a huge effort was deployed at all ends to build a system that can handle the data intake.
- Hard and software solutions have been built to be **scalable** and flexible to allow to **adapt to change**.
- The goal of ALMA is to help the scientists wherever possible from proposal preparation over the **science-grade data products** and data-reduction to archival research and to work towards a great end-to-end **user-experience**