VLBI Data Systems

Alan R. Whitney
MIT Haystack Observatory

2 March 2013
VLBI School
Helsinki, Finland

Only options to improve VLBI sensitivity are . . .

- Bigger antennas, but cost tends to go as $\sim D^{2.7}$
- Quieter receivers, but many receivers are already approaching quantum noise limits or are dominated by atmospheric noise
- Wider observing bandwidth
  - For most VLBI observations, sensitivity increases as square root of observed bandwidth
  - Increasing BW is usually the most cost-effective way to increase sensitivity
  - As a result, VLBI has always pushed data-recording technology to its limits!

What data are actually recorded?

Answer: It is nothing more than precisely timed samples of pure Gaussian noise!

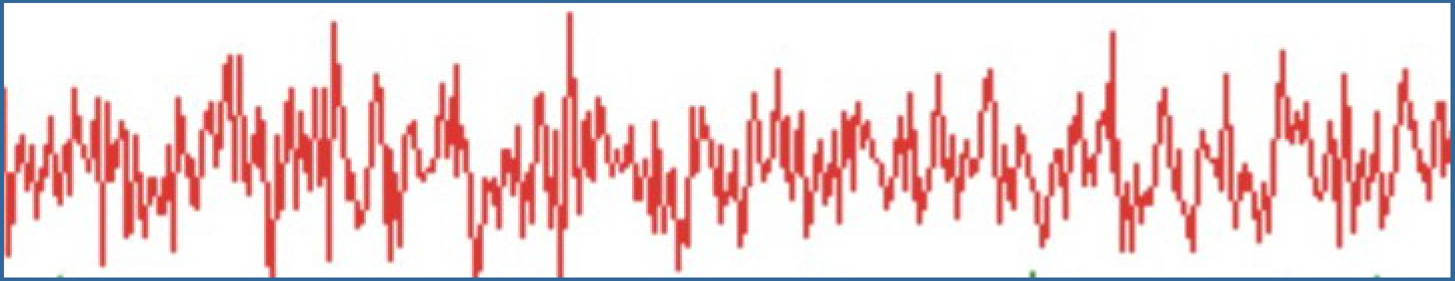

Interesting fact: Normally, the voltage signal is sampled with only 1 or 2 bits/sample

Interesting fact: VLBI data are essentially incompressible (particularly continuum VLBI)

Important fact:

If a small amount of data are lost, it's usually no big deal!

Why only 1 or 2 bits/sample?

- **In 1960's John Van Vleck of Univ. of Wisconsin showed that:**
  - The spectrum of a Gaussian-statistics bandwidth-limited signal may be completely reconstructed by measuring only the sign of the voltage at each Nyquist sampling point!!!
- **For VLBI:**

  If sampling at ∞ bits/sample at a given BW produces an SNR of 1.0, then:
  - Sampling at 1 bit/sample produces an SNR of ~0.63 compared to ideal analog of 1.0
  - Sampling at 2 bits/sample produces an SNR of ~0.87 compared to ideal analog of 1.0
  - Recall: SNR increases as $\sqrt{BW}$

This is why!

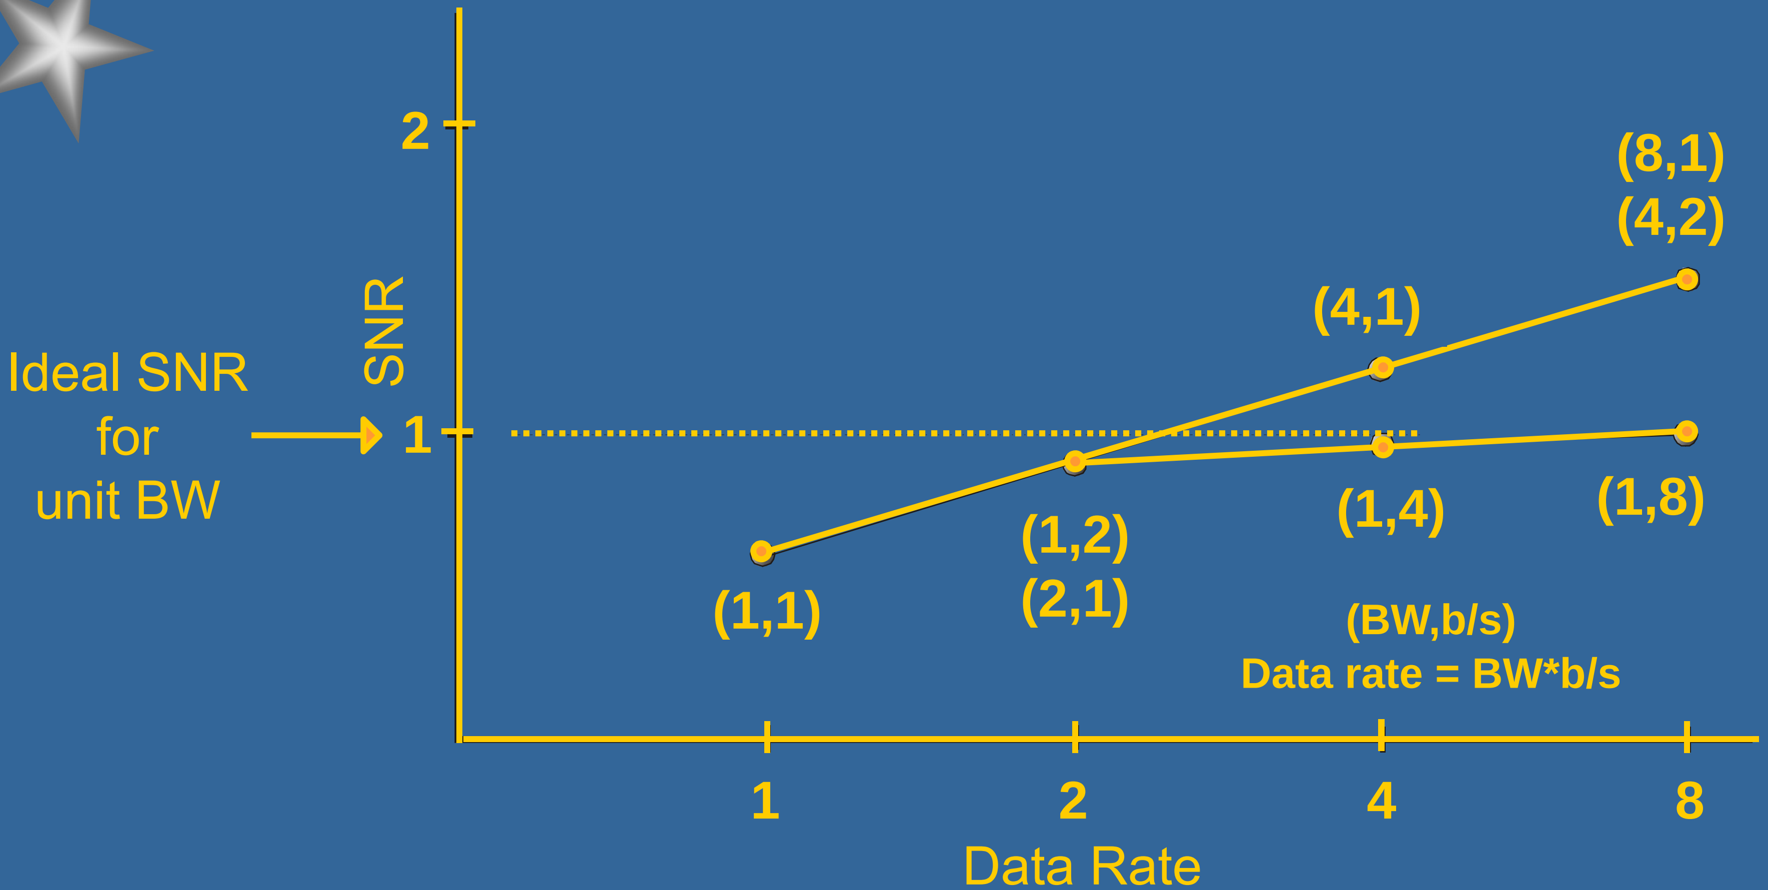

Note that SNR goes up faster by increasing BW than by increasing #bits/sample

Conclusion: To maximize SNR when data-rate is constrained, it is best to increase the BW !

Example of sampling a waveform at 2 bits/sample

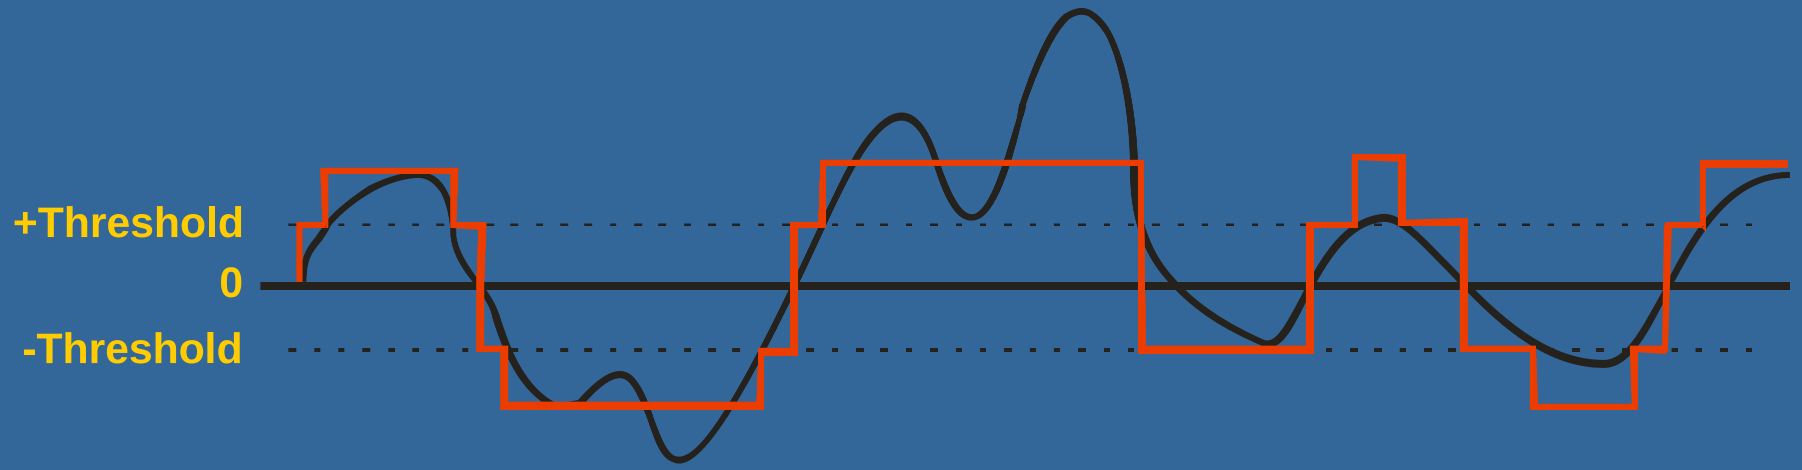

Cross-correlation of weak signals in noise

- Let s(t) be a weak astronomical signal, and $n_1(t)$ and $n_2(t)$ be noise signals at sites 1 & 2

Receiver 1 noise $n_1(t)$

Receiver 2 noise $n_2(t)$

Signal s(t)

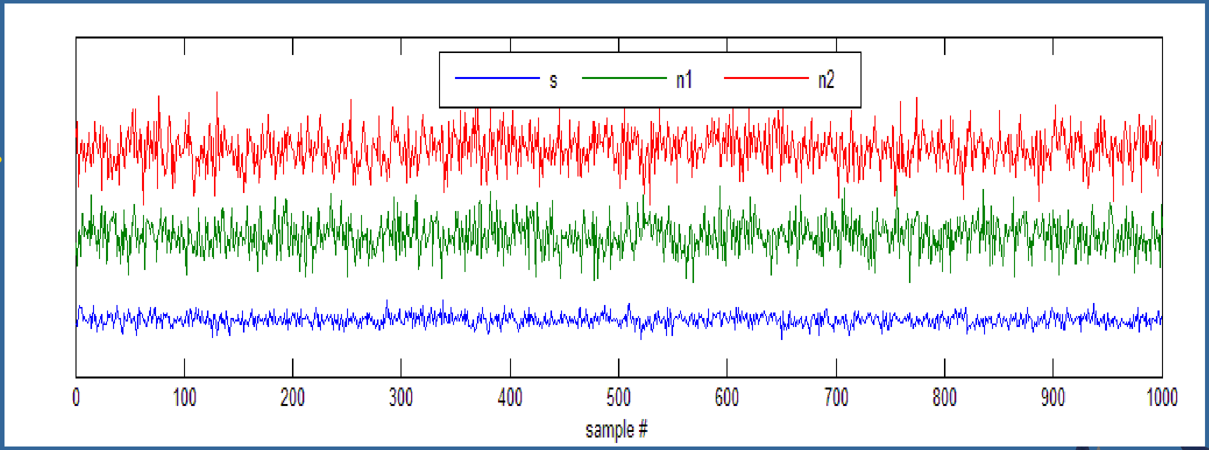

Cross-correlation of weak signals (cont'd)

Product of signals is:

$$(s + n_1)\,(s + n_2) = s^2 + n_1 s + n_2 s + n_1 n_2$$

In actuality, life is more complicated due to Earth rotation:

- Time-of-arrival difference continually changes
- Differential Doppler shift continually changes

Correlation components

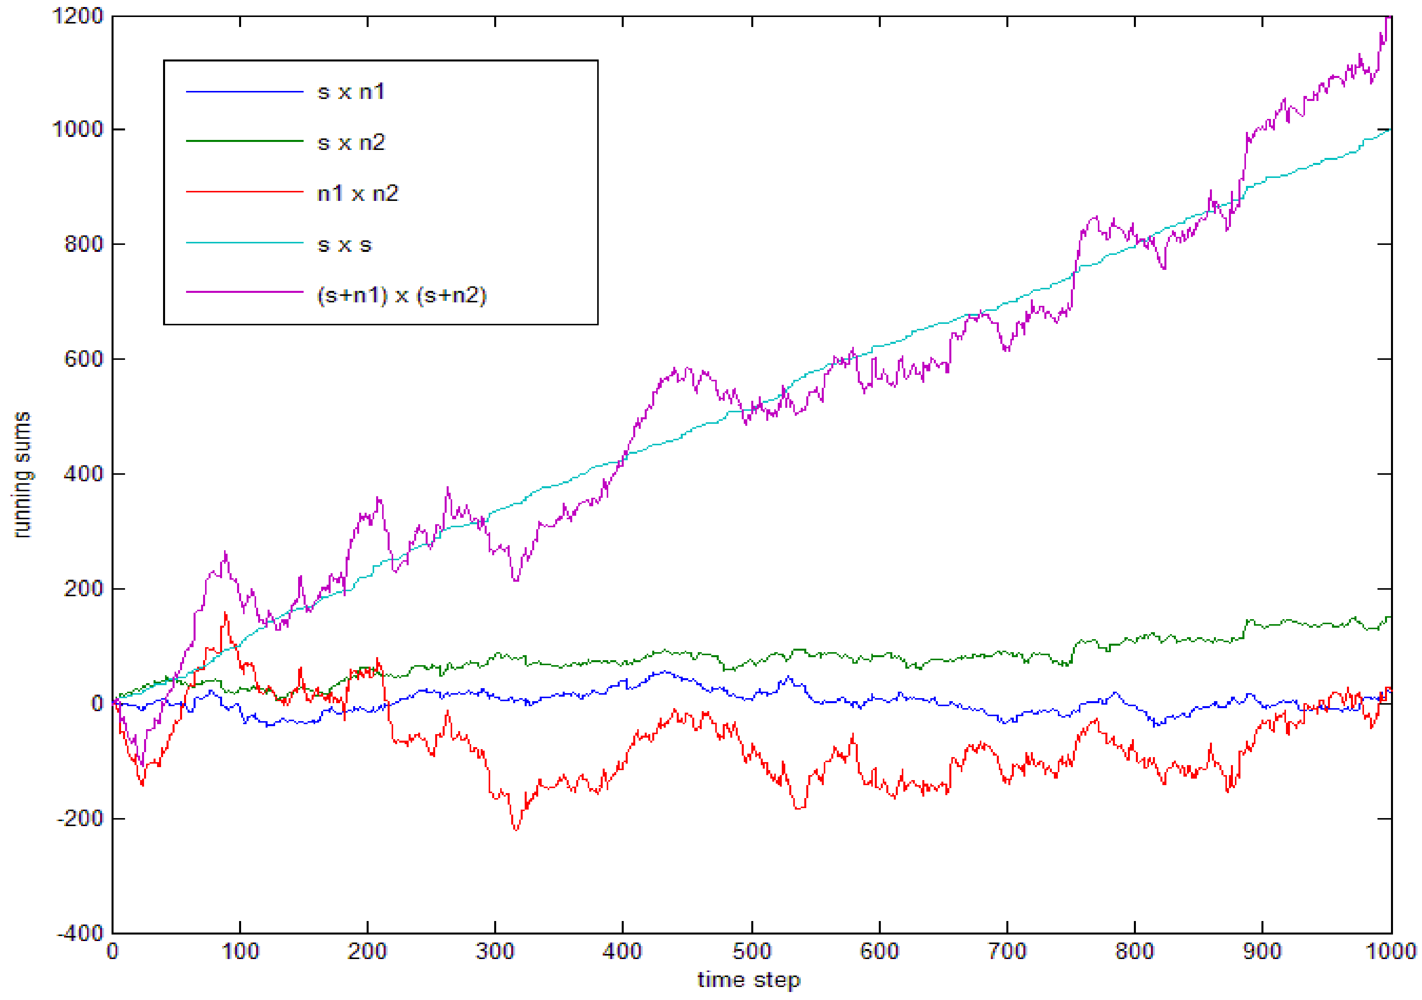

Data rates for geodetic-VLBI

- geodetic-VLBI data rates are dictated by
  - Small antennas (12m class)
    - Antennas must be able to move very quickly around sky
  - Weak sources
    - Sources need to be ~uniformly distributed in the sky and have simple or no structure; severely constrains candidates
  - Short observations
    - VLBI2010 observations will be 30-60 secs each
    - antenna must move around sky quickly to map temporal atmospheric flucuations
    - most of observation period is spent moving antenna from source-to-source

All these factors conspire to dictate very high data rates (both instantaneous and average)

VLBI data-rate drivers

- Geodetic VLBI (VLBI2010 project)
  - Near term: 16Gbps burst capture; 4Gbps sustained
  - Long term: ~32/64Gbps burst capture; ~8/16Gbps sustained
- mm-VLBI
  - Can use 16Gbps as soon as available
  - Phased-ALMA will be able to support 64Gbps

VLBI Data Rates and Volume are not for the faint of heart!

- Experiments at 1-4 Gbps/station, 4 to 20 stations
  - ~5-40 TB/station/day
  - Global 10-station experiment @ 4 Gbps/station up to ~400 TB/day
  - Single 10-day experiment can produce up to ~4 PB
- Higher data rates (8-32 Gbps) are already on the horizon; higher data rates → more sensitivity
- Available disk supply can support only few days of observations at these rates
- All pairwise telescope combinations must be cross-correlated

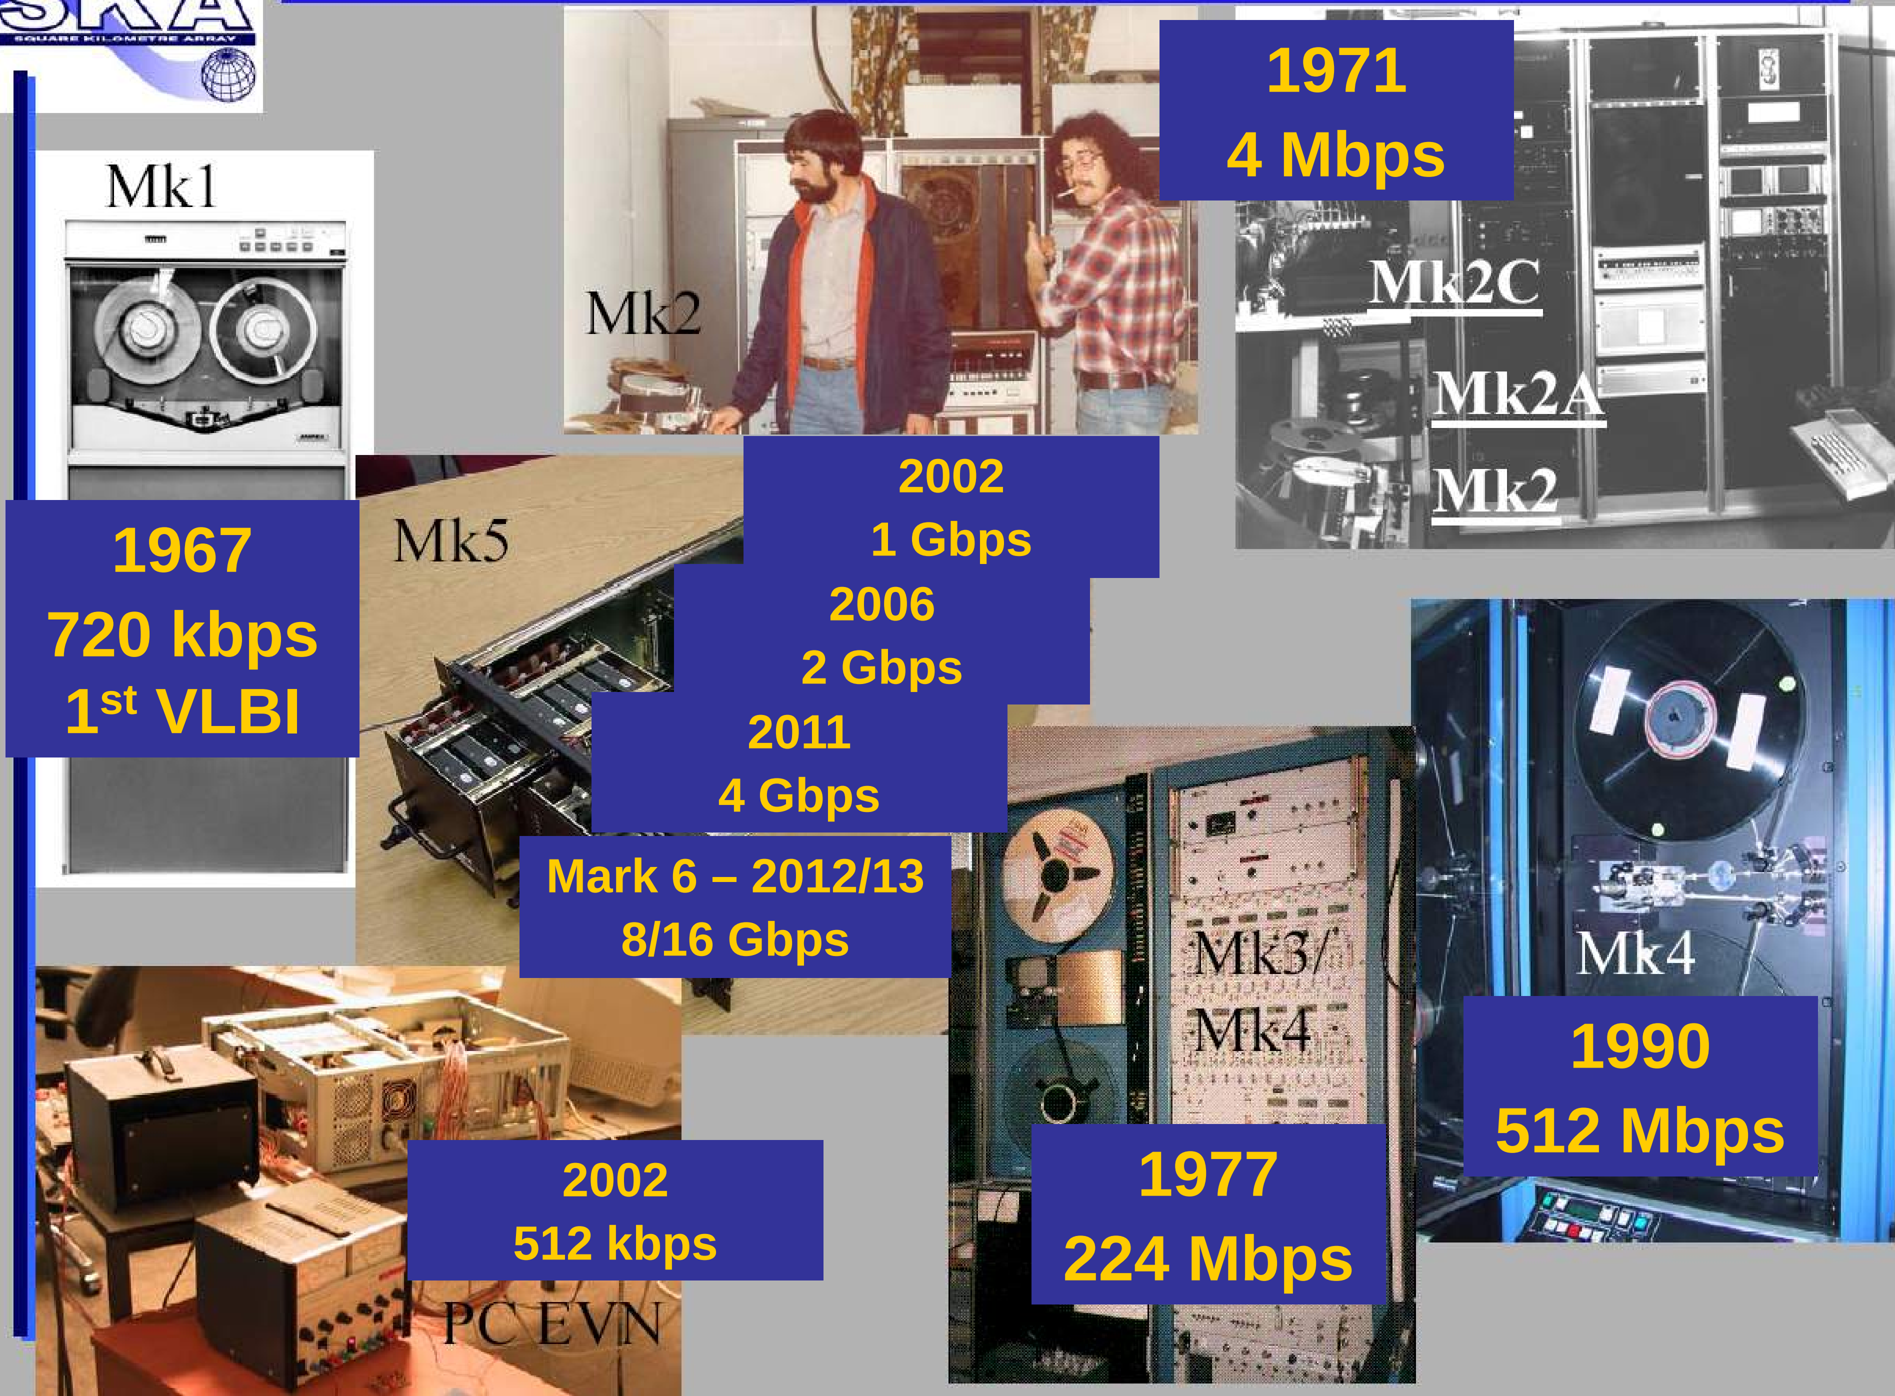
Mk1
1967
720 kbps
1st VLBI
Mk2
1971
4 Mbps
Mk2C
Mk2A
Mk2
Mk5
2002
1 Gbps
2006
2 Gbps
2011
4 Gbps
Mark 6 – 2012/13
8/16 Gbps
Mk3/
Mk4
1977
224 Mbps
Mk4
1990
512 Mbps
2002
512 kbps
PC EVN

VLBI Interfaces

- Mark 5/VLBA parallel bit-streams
  - Track-based bit-streams designed for tape
  - Rather *ad hoc*, but widely used by tape-based systems and some early disk systems (Mark 5A, for example)
- VLBI Standard Interface-Hardware (VSI-H)
  - ratified by community vote at international VLBI meeting
  - developed by appointed VSI Task Force in 1999
  - won Japanese Ministry award and adoption by some other disciplines as well; re-christened as 'Versatile Streaming Interface'
  - used by several VLBI data systems, including Mark 5B/5B+, PC-EVN, K5
- 10 Gigabit Ethernet
  - now becoming standard interface for nearly all systems
  - Largely data-format independent

Design considerations for VLBI recording systems

- Support for standard VLBI data interface(s)
- Removable modular data-storage units
  - Needed for convenient shipping of data
  - For disk modules: easy replacement of individual disks within modules
  - Accommodation for spare ‘ready’ set of media for automatic changeover when active media are filled (‘ping-pong’ type of operation)
- Resilient real-time recording
  - Design to avoid data-flow slowdown in face of individual slow recording elements (i.e. disks)
  - Occasionally drop data segments if necessary to keep up with data flow (few percent data-loss usually not crippling)
- Support for e-VLBI if possible

Example: Mark 5 Data Acquisition Systems

(Mark 5A/B/B+/C all look the same)

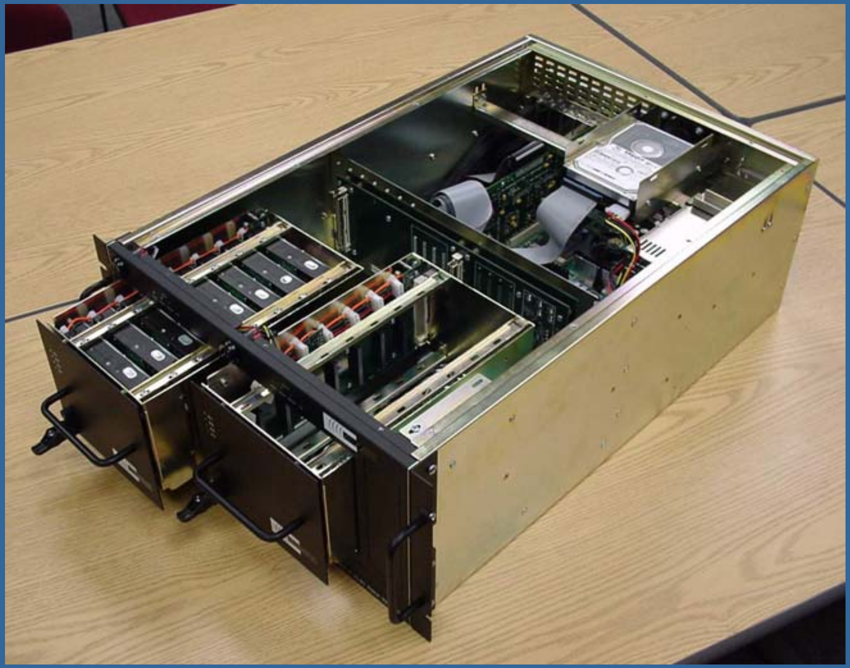

	Year introduced	Record rate (Mbps)	Interface	Cost (USk$)	#deployed
Mark 5A	**2002**	**1024**	**Mk4/VLBA**	**21**	**~130**
Mark 5B	**2005**	**1024**	**VSI-H**	**22**	**~40**
Mark 5B+	**2006**	**2048**	**VSI-H**	**23**	**~30**
Mark 5C	**2011/12**	**4096**	**10GigE**	**21**	**~30**

Mark 5 includes a significant amount of proprietary technology

Other disk-based VLBI data systems

- PC-EVN (from Metsahovi)
  - Custom-designed plug-in card for standard PC
  - 512Mbps sustainable record rate
  - Recording to disk array
  - VSI-H data interface
- K5 (NICT)
  - VSI-H data interface
  - 1Gbps sustainable record rate to disk array

VLBI Data Formats

- Data formats were quite a varied lot until Mark4/VLBA tape systems became prevalent
  - Track-based formats designed specially for unreliabilities of tape
  - Have largely been superseded by more modern data formats designed for disk-based systems
- Mark 5B data format
  - designed to complement VSI-H hardware interface
  - won Japanese Ministry award and adoption by some other disciplines as well; re-christened as ‘Versatile Streaming Interface’
  - used by several VLBI data systems designed in U.S., Japan, China and Russia; still in use
- VLBI Data Interchange Format (VDIF)
  - developed by an appointed VLBI Task Force and ratified by the VLBI community in 2009
  - designed to accommodate a wide variety of data-recording/transmission needs
  - virtually all new systems being designed to accommodate VDIF

Disk-based data formats

- Mark 5B data format
  - designed to complement VSI-H hardware interface
  - used by at least two disk-based systems (Mark 5B/
  - won Japanese Ministry award and adoption by some other disciplines as well; re-christened as ‘Versatile Streaming Interface’
  - used by several VLBI data systems designed in U.S., Japan, China and Russia (some still in use)
- VLBI Data Interchange Format (VDIF)
  - developed by an appointed VLBI Task Force and ratified by the VLBI community in 2009
  - designed to accommodate a wide variety of data-recording/transmission needs
  - virtually all new systems being designed to accommodate VDIF

VDIF
(VLBI Data Interchange Format)

- Standardized format for raw time-sampled VLBI data
- Compatible with both VLBI data-recording systems and e-VLBI data transmission
- Highly flexible to accommodate a large variety of channel and frequency configurations, including mixed sample-rate data; supports rates to ~100Gbps
- VDIF being implemented for all new VLBI2010 systems
- Accompanying VLBI Transport Protocol (VTP) specifies e-VLBI data-transmission protocol for VDIF-formatted data stream

For details: www.vlbi.org

VDIF Assumptions

- Data are assumed to be one or more time series of uniformly time sampled data
- Each time series may have own sample rate, bits/sample and station of origin

VDIF Requirements

- Data are self-identifying wrt time tag and data identification
- Data may be discontinuous in time (e.g. pulsar data)
- Data may be single-channel or multi-channel
- Data may be single bit or multi-bit samples
- Data can be decoded with no external input
- Number of channels should be arbitrary (i.e. not confined to 2^n)
- Suitable for on-wire transfer as well as disk file storage
- Support data rates up to at least 100Gbps

VDIF Structure

- A ‘Data Stream’ is organized into self-identifying ‘Data Threads’
- Each Data Thread contains of a serial set of ‘Data Frames’
- Data Frame length may be chosen by user, but must be same for all threads
- Data Frame may contain single-channel or multi-channel data
- Each Data Thread may have different number of channels, sample rate, bits/sample
- For Ethernet connection, VDIF frame length must be <~9kB

VDIF Structure

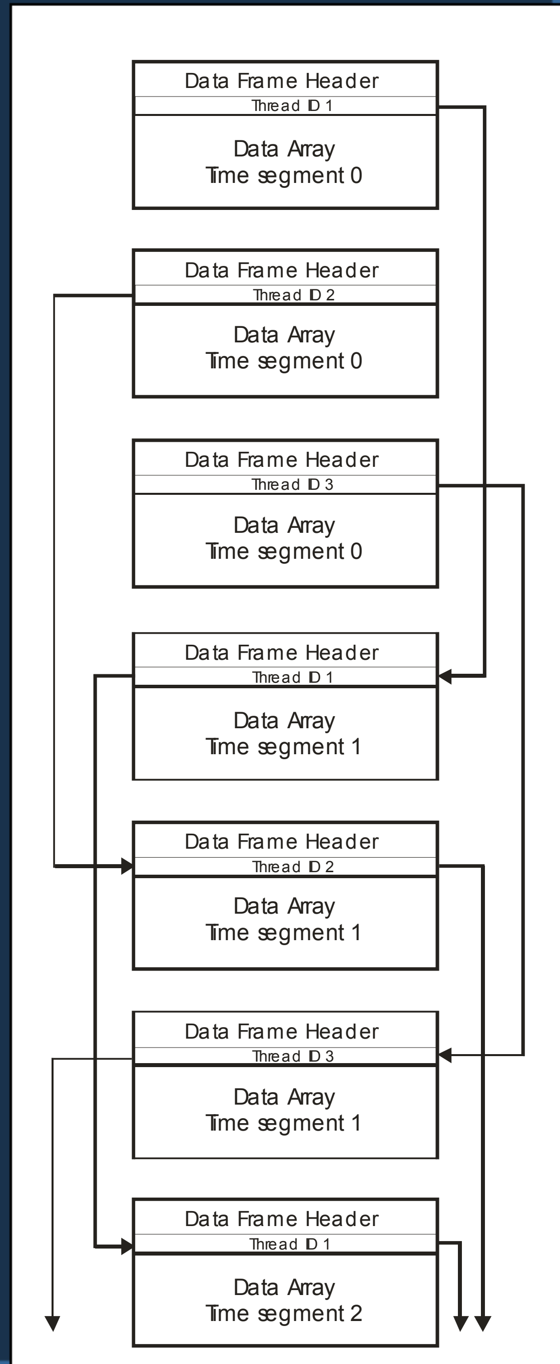

VDIF Data Frame Header Content

- Time (seconds since 00UT of defined 6-month period during which there can be no leap-seconds)
- Frame # within second
- Stream ID
- Station ID (2-char ASCII code)
- ‘Data-invalid’ marker
- #channels
- Bits/sample
- Data Array length
- VDIF version #
- Optional user-defined 16-byte extension

Data Frame

- Each Data Frame has 16/32 byte header followed by a Data Array of user-specified length
- Data Frame length for a single Data Thread is fixed for a particular scan
- Data Frame length must be a multiple of 8 bytes
- Data Frames per second must be an integer
- Data Frame may not span a second boundary

Typical Usage

- Single Thread of multi-channel Data Frames
  - Multiple channels in a single Data Stream
  - Avoids ‘corner turning’ requirement
- Multiple Threads of single-channel Data Frames
  - Multiple channels using multiple threads
  - Preferred for new equipment and applications

VDIF-compatible systems

- Building a VDIF-compatible system not hard – just capture and record incoming VDIF-formatted Ethernet packets on a standard PC
- But -- building a VDIF-compatible system for VLBI field use is a bit more complicated
  - need to meet all the design criteria we talked about earlier

Mark 6 – a VDIF compatible system

- 8Gbps demonstrated sustained record and playback capability; goal of 16Gbps sustainable
- 4 x 10GigE input ports
- Based on inexpensive high-performance COTS hardware
- Linux OS w/open-source software
- Resilient File System to manage slow and failed disks
- e-VLBI support
- Preserve as much investment in existing Mark 5 systems and disk libraries as possible
- Mark 6 collaboration:
  - Haystack Observatory – all software and software support; hardware specification
  - NASA/GSFC High-End Network Computing group – consultation on high-performance COTS
  - Conduant Corp –Mark 6 disk module, disk-module power management

Basic Mark 6 System

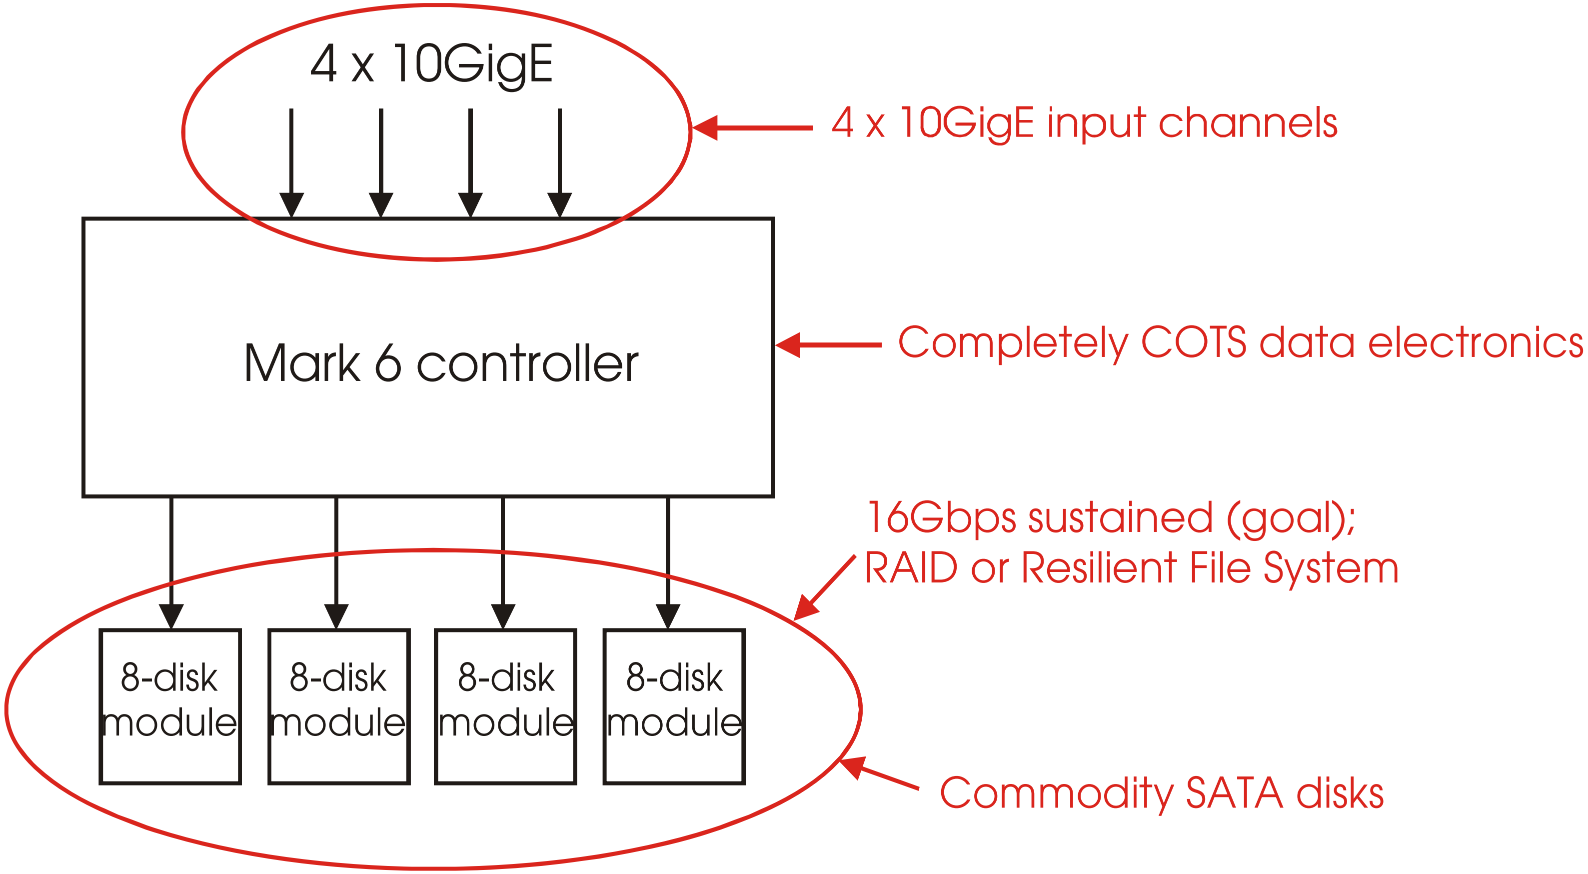

Prototype Mark 6 hardware

16 Gbps VLBI demonstration with Mark 6

24 October 2011

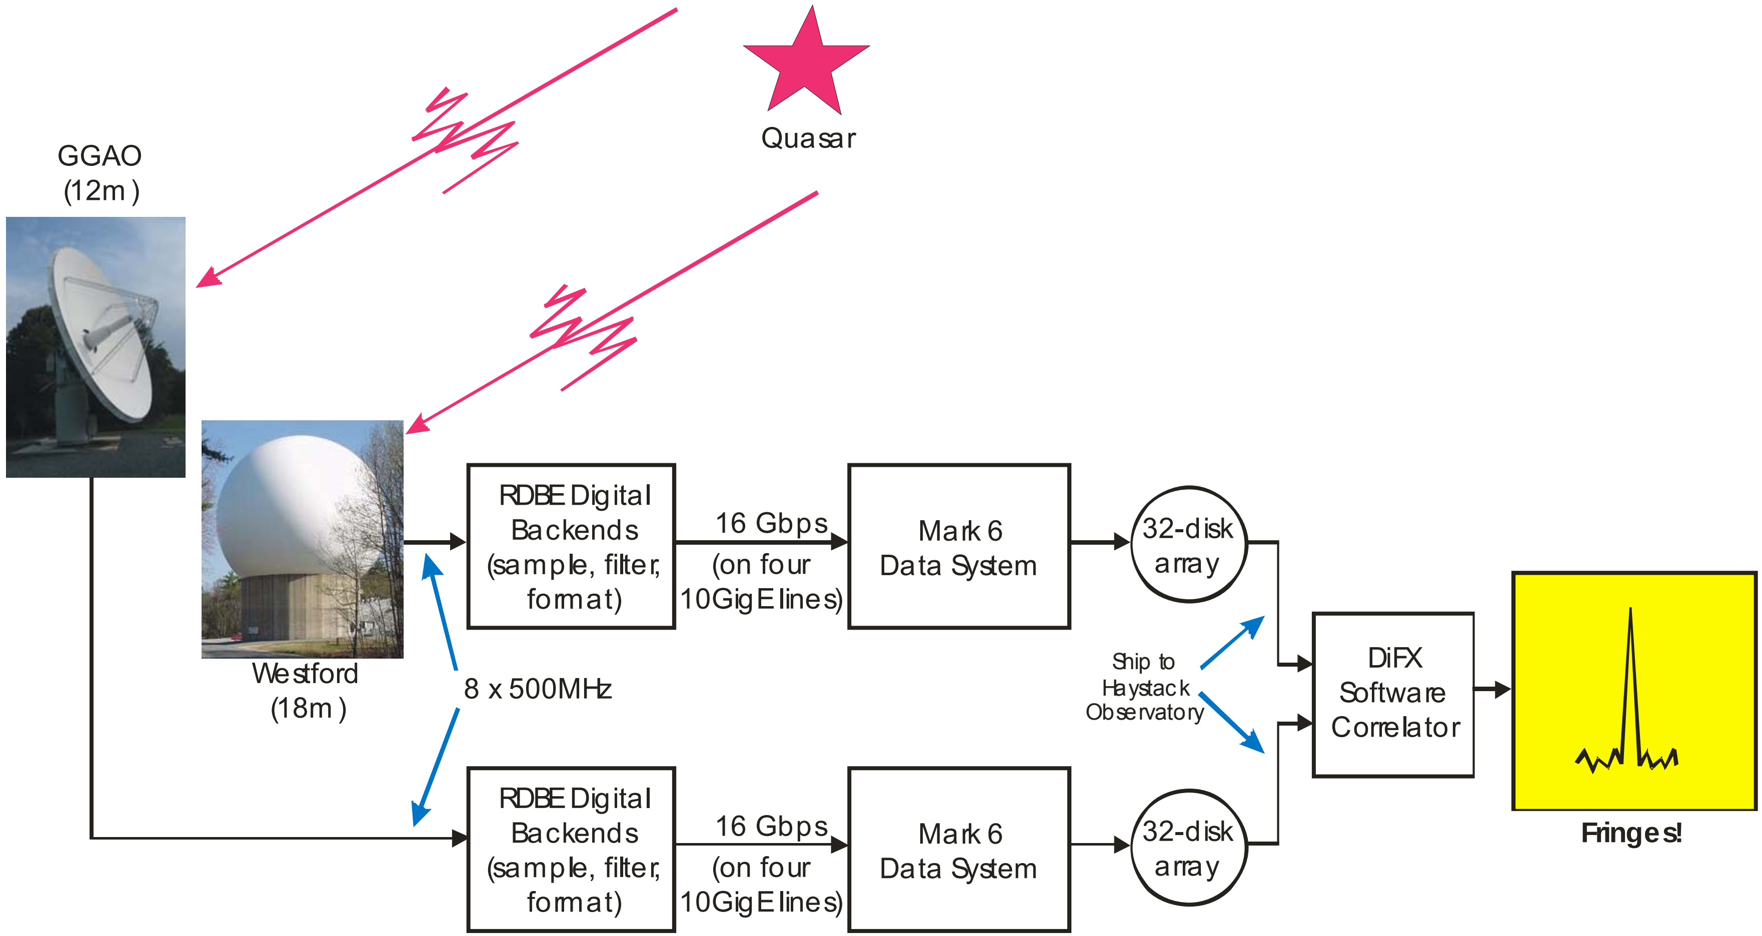

Correlation results (single 500MHz channel)

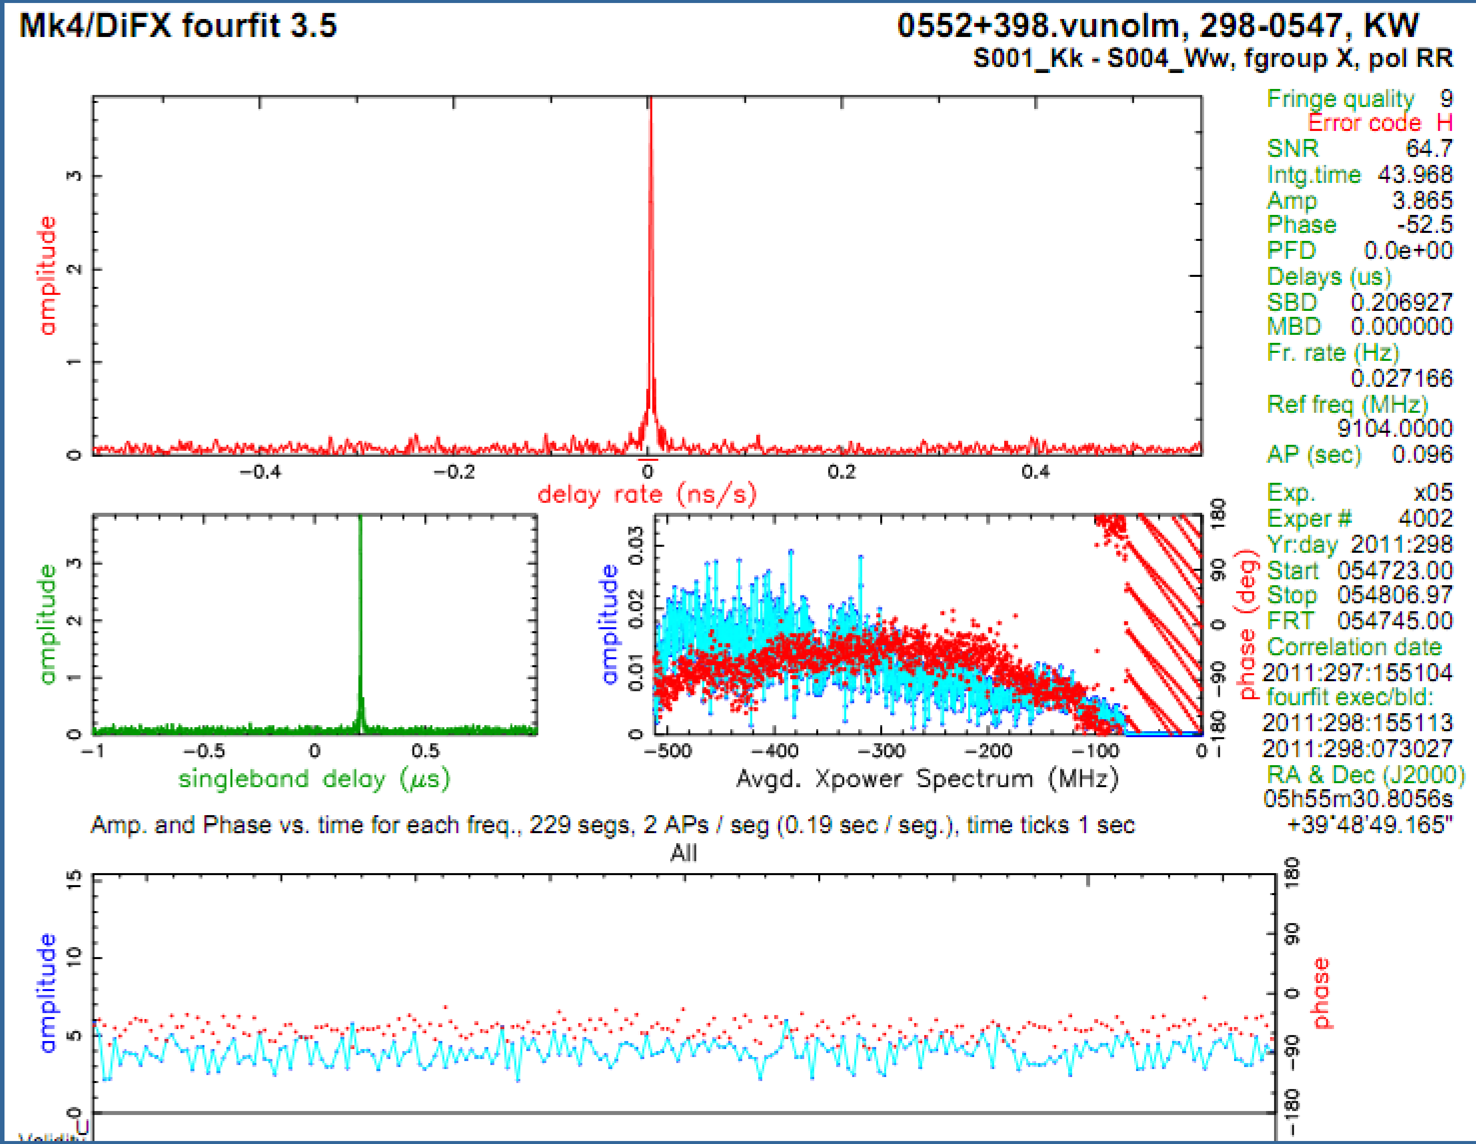

Recording-rate capability vs time

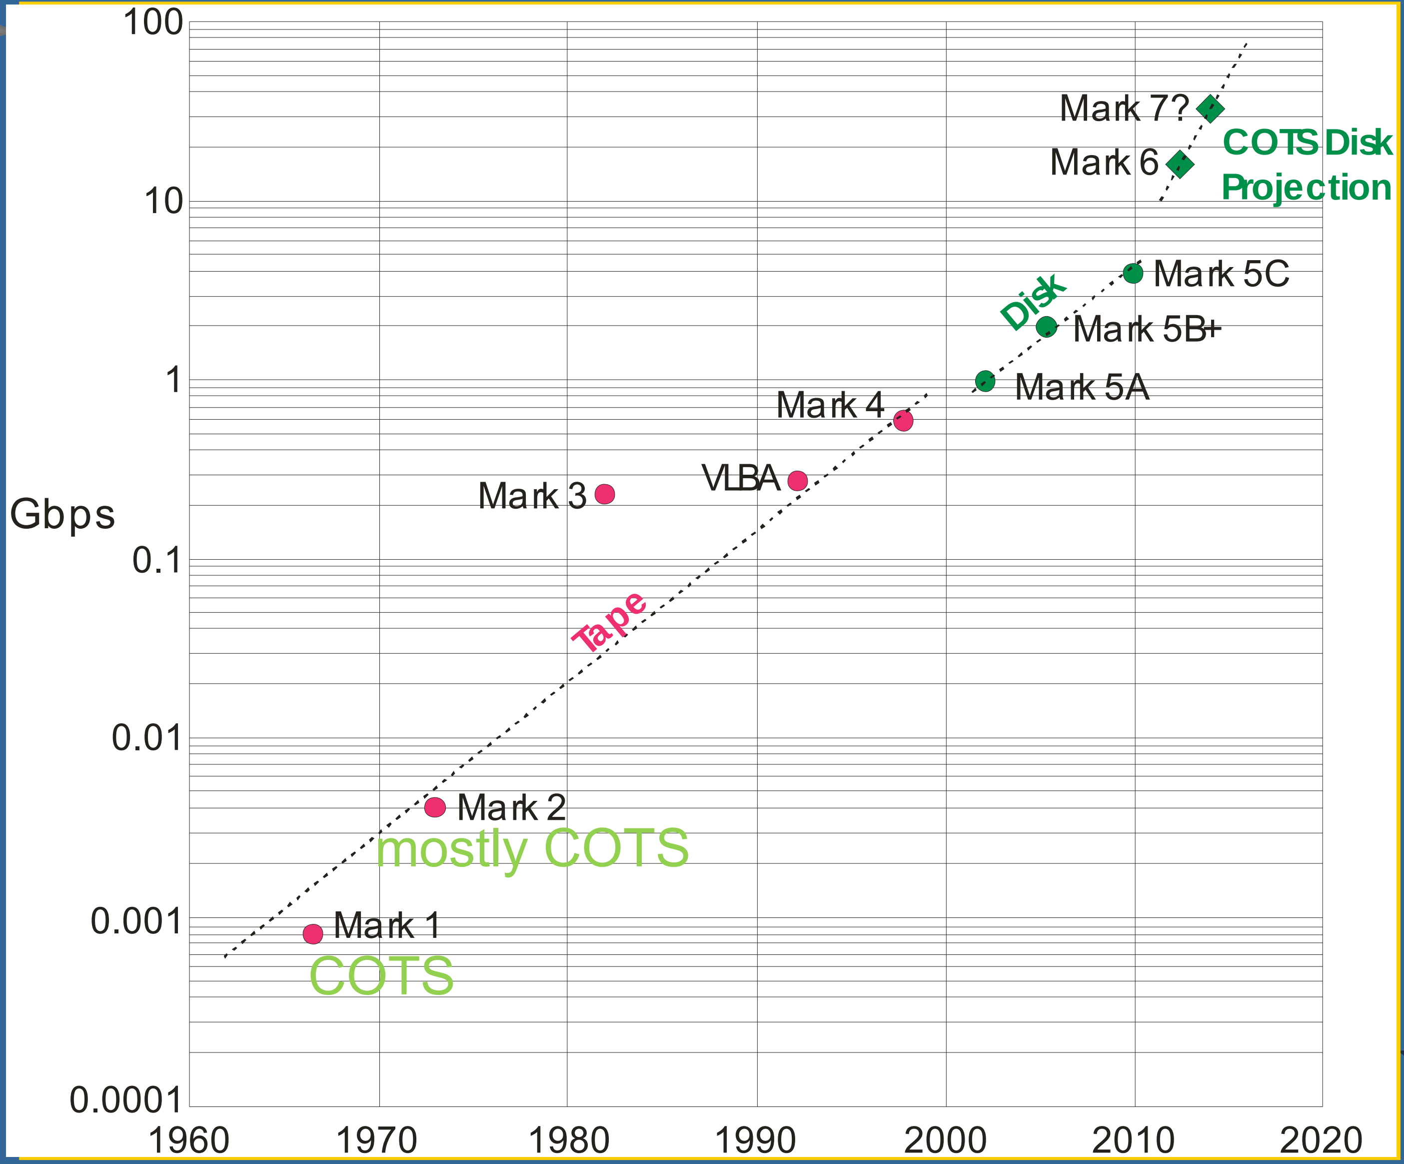

Backend-bandwidth cost vs. time

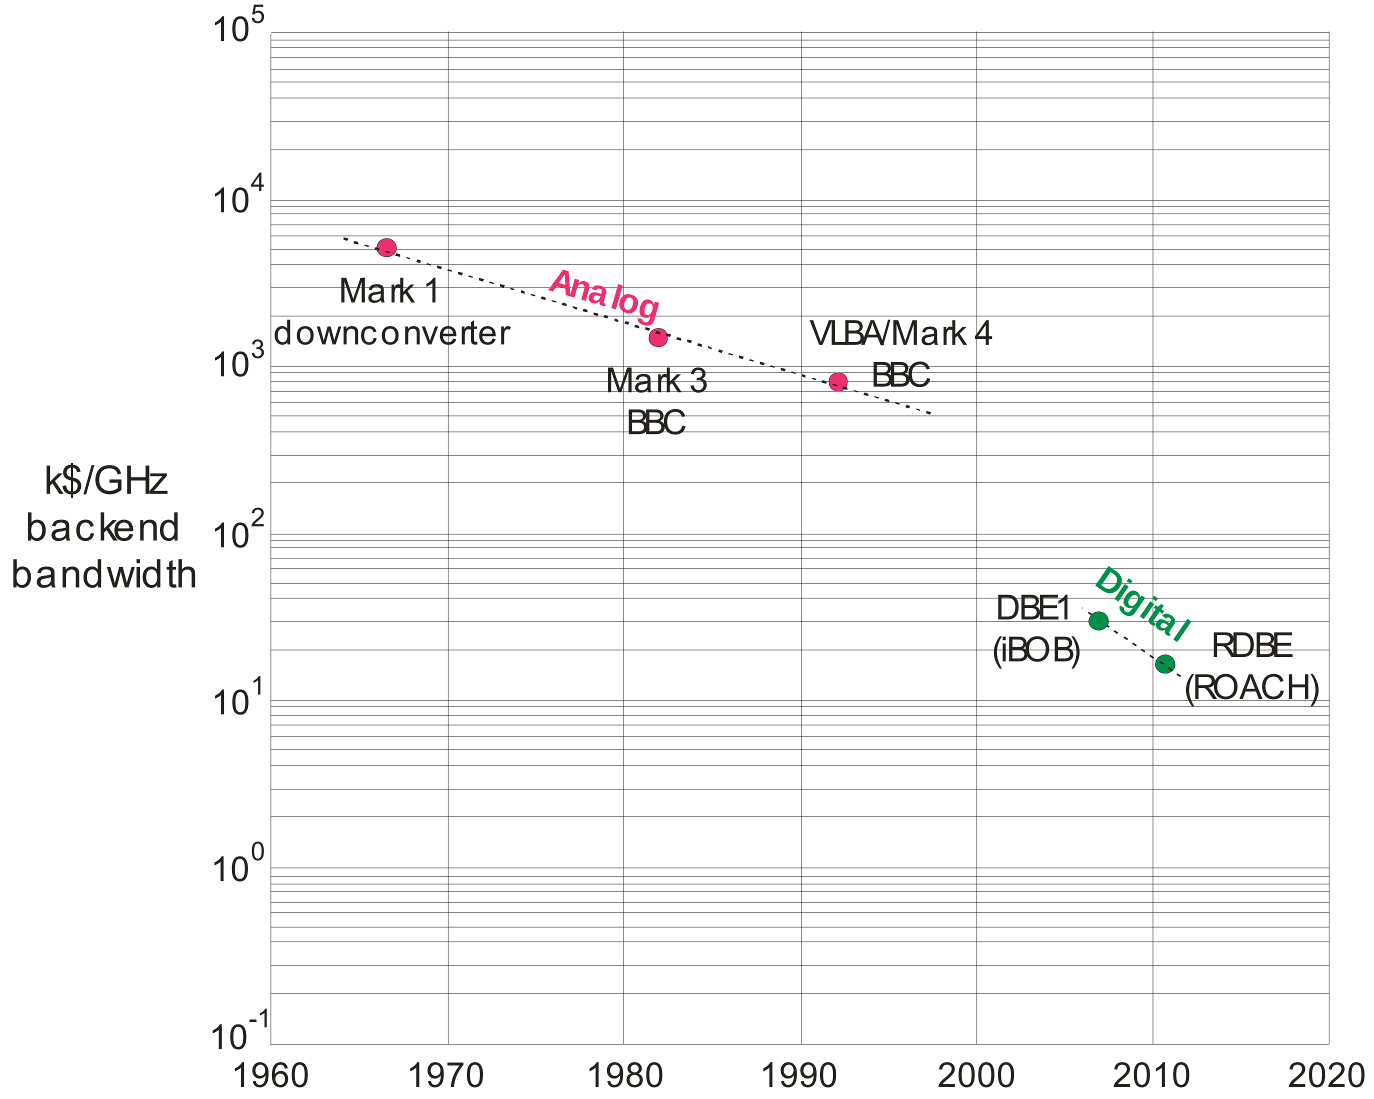

Recording-rate cost vs. time

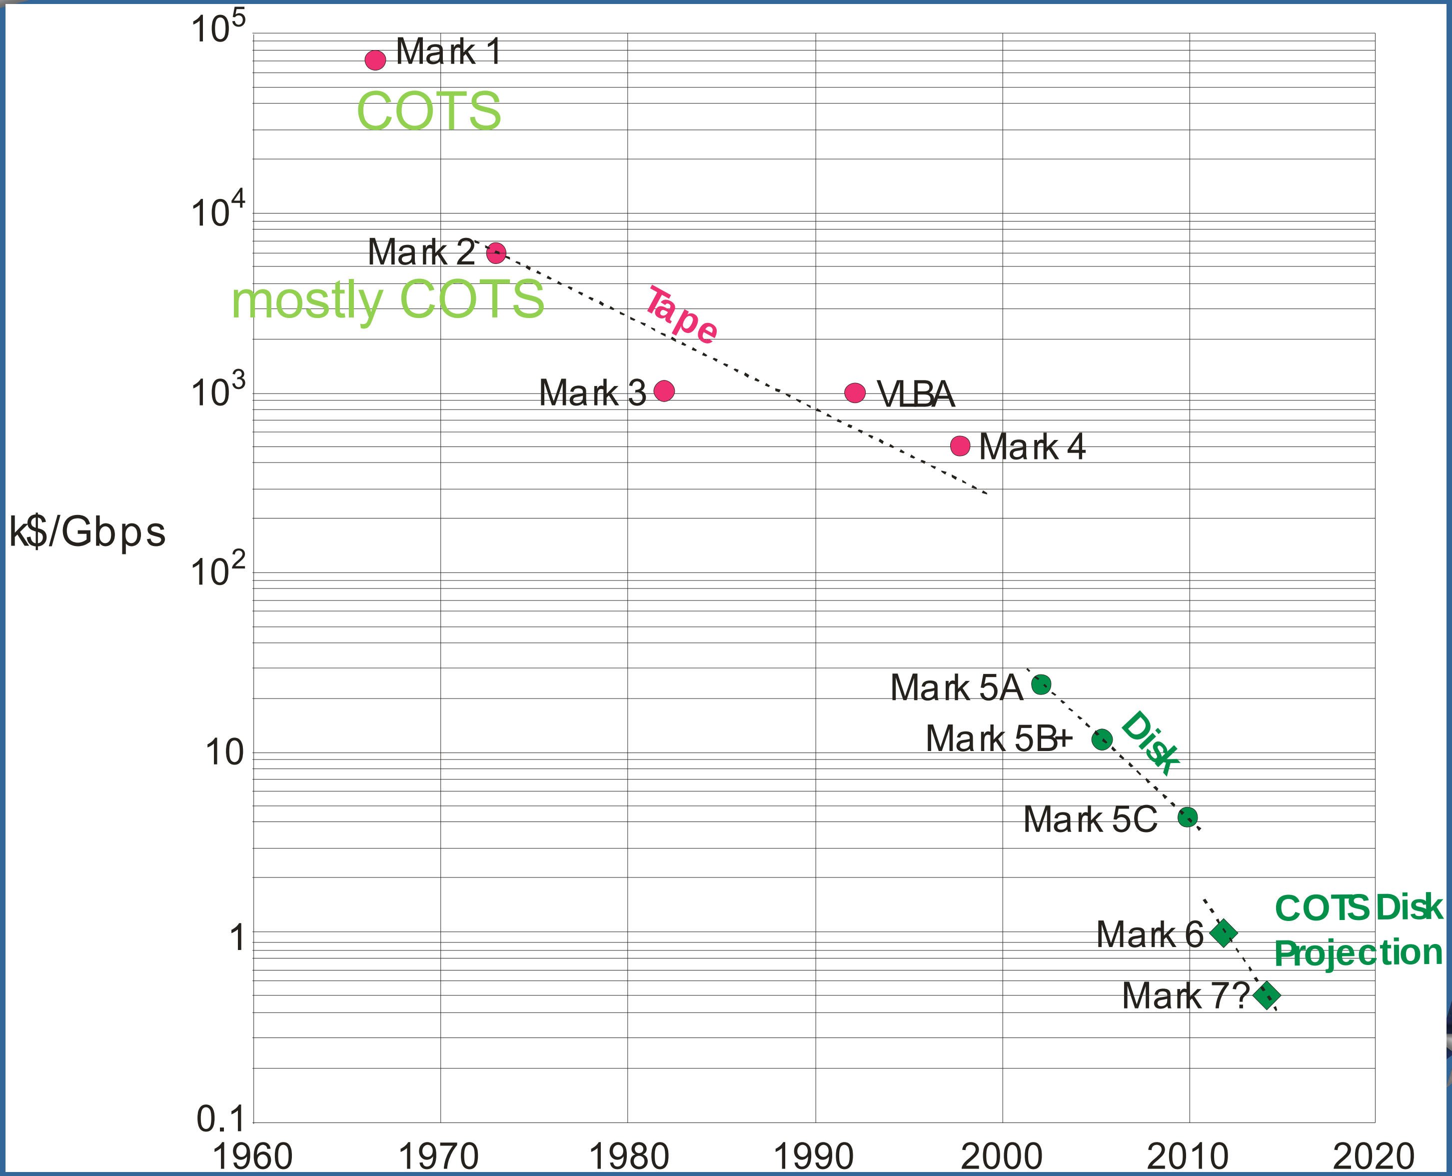

Tape vs. Disc Price Comparison

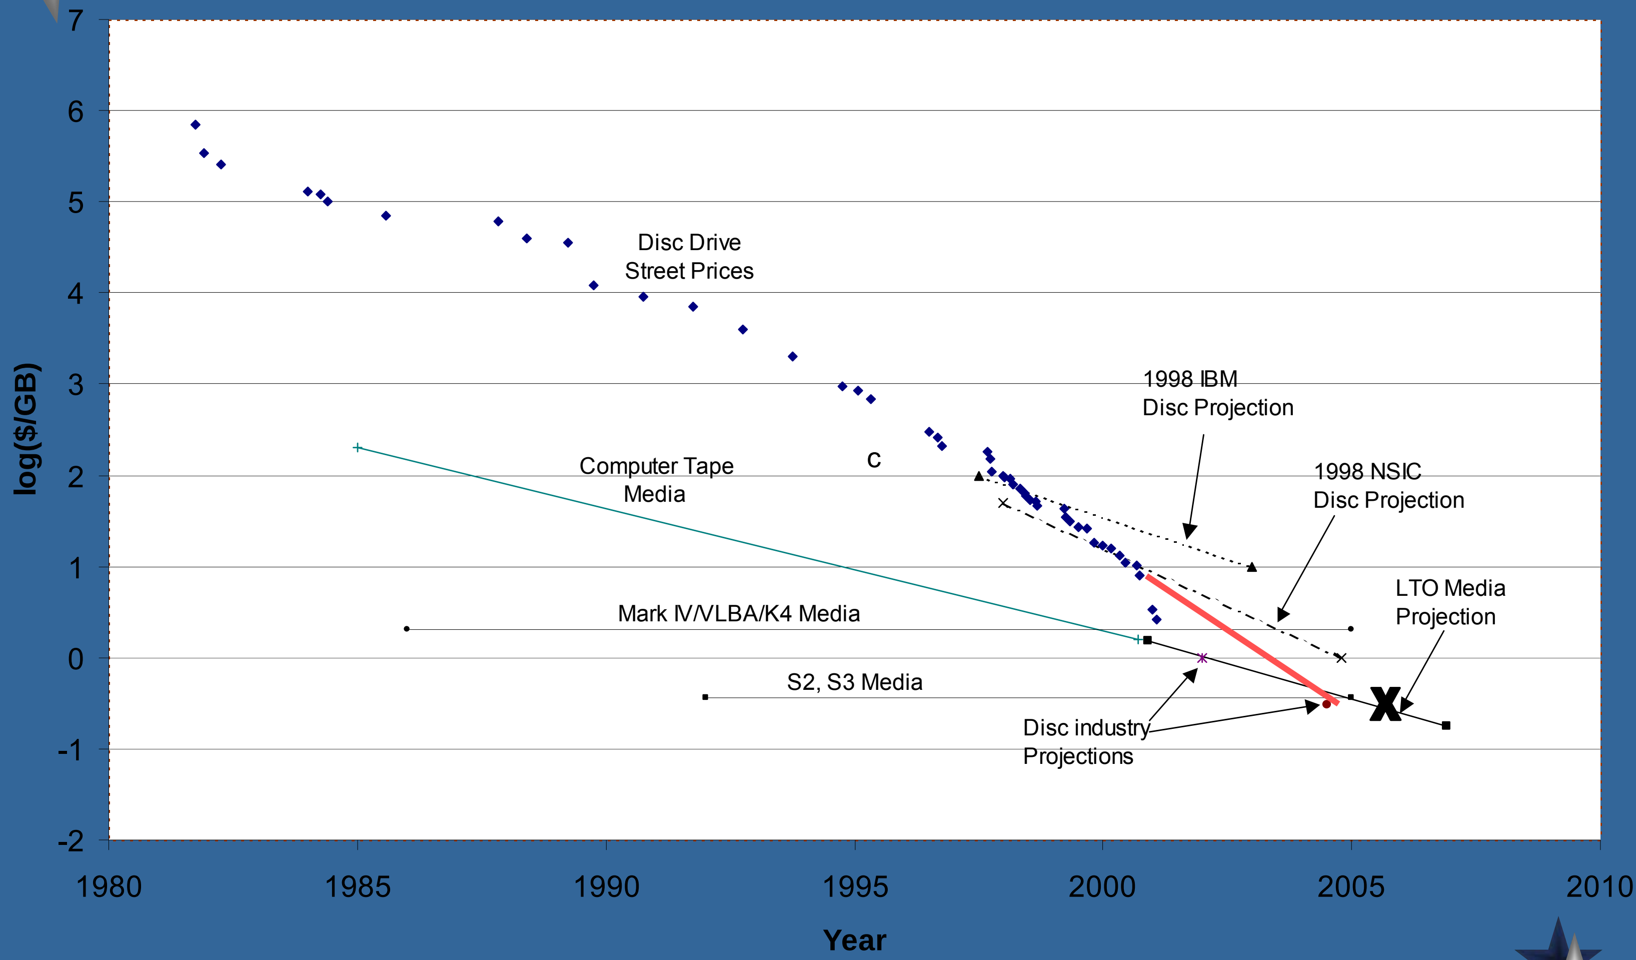

Enter 'e-VLBI': Electronic Transmission of VLBI Data

Of course, not a new idea, but only recently becoming practical and affordable

- **1977** – Canadian's used a satellite to transmit data in real-time from Green Bank, WV to Algonquin, Canada at 20 Mbps (pretty impressive for the time!)
- **1979** – Haystack developed near-real time correlation using data transmitted at 1200 bps over POTS using computer modems
- **Mid-1990's** – Japanese developed dedicated 4-station network around Tokyo operating at 256 Mbps over dedicated fiber-optic links
- **2000's** – e-VLBI is slowly becoming a growing force as stations are connected; Japan and Europe are the leaders
- **2012** – e-VLBI is standard practice for increasing fraction of European astronomy and geodesy; limited e-VLBI, mostly non-real-time gaining groung elsewhere
- **VLBI2010** – e-VLBI a necessity if VLBI2010 is to meet goals of rapid turnaround!

Advantages of e-VLBI

Increased Reliability

- remote performance monitor & control capability in near real-time
- avoid unexpected media-shipping problems

Rapid turnaround of time-sensitive observations

- particularly important for UT1 observations

Lower Cost

- Enables full automatic operation
  Eliminates manual handling and shipping of storage media
- Elimination or reduction of recording-media pool (millions of $’s!)
- VLBI data systems such as Mark 6 can be dual-use
  i.e. suitable for both e-VLBI and recording

But there are big challenges!

- **Network bottlenecks are often well below advertised rates**
- **Performance of transport protocols**

    -Untuned TCP stacks, fundamental limits of regular TCP

- **Throughput limitations of COTS hardware**

    -Network interfaces

    -Disk-I/O - Network

- **Complexity of e-VLBI experiments**

    -e-VLBI experiments currently require significant network expertise

- **Time-varying nature of network**
- **‘Last-mile’ connectivity to telescopes**

    -Most telescopes are deliberately placed in remote areas

Europe is moving ahead with EXPRes/NEXPReS projects

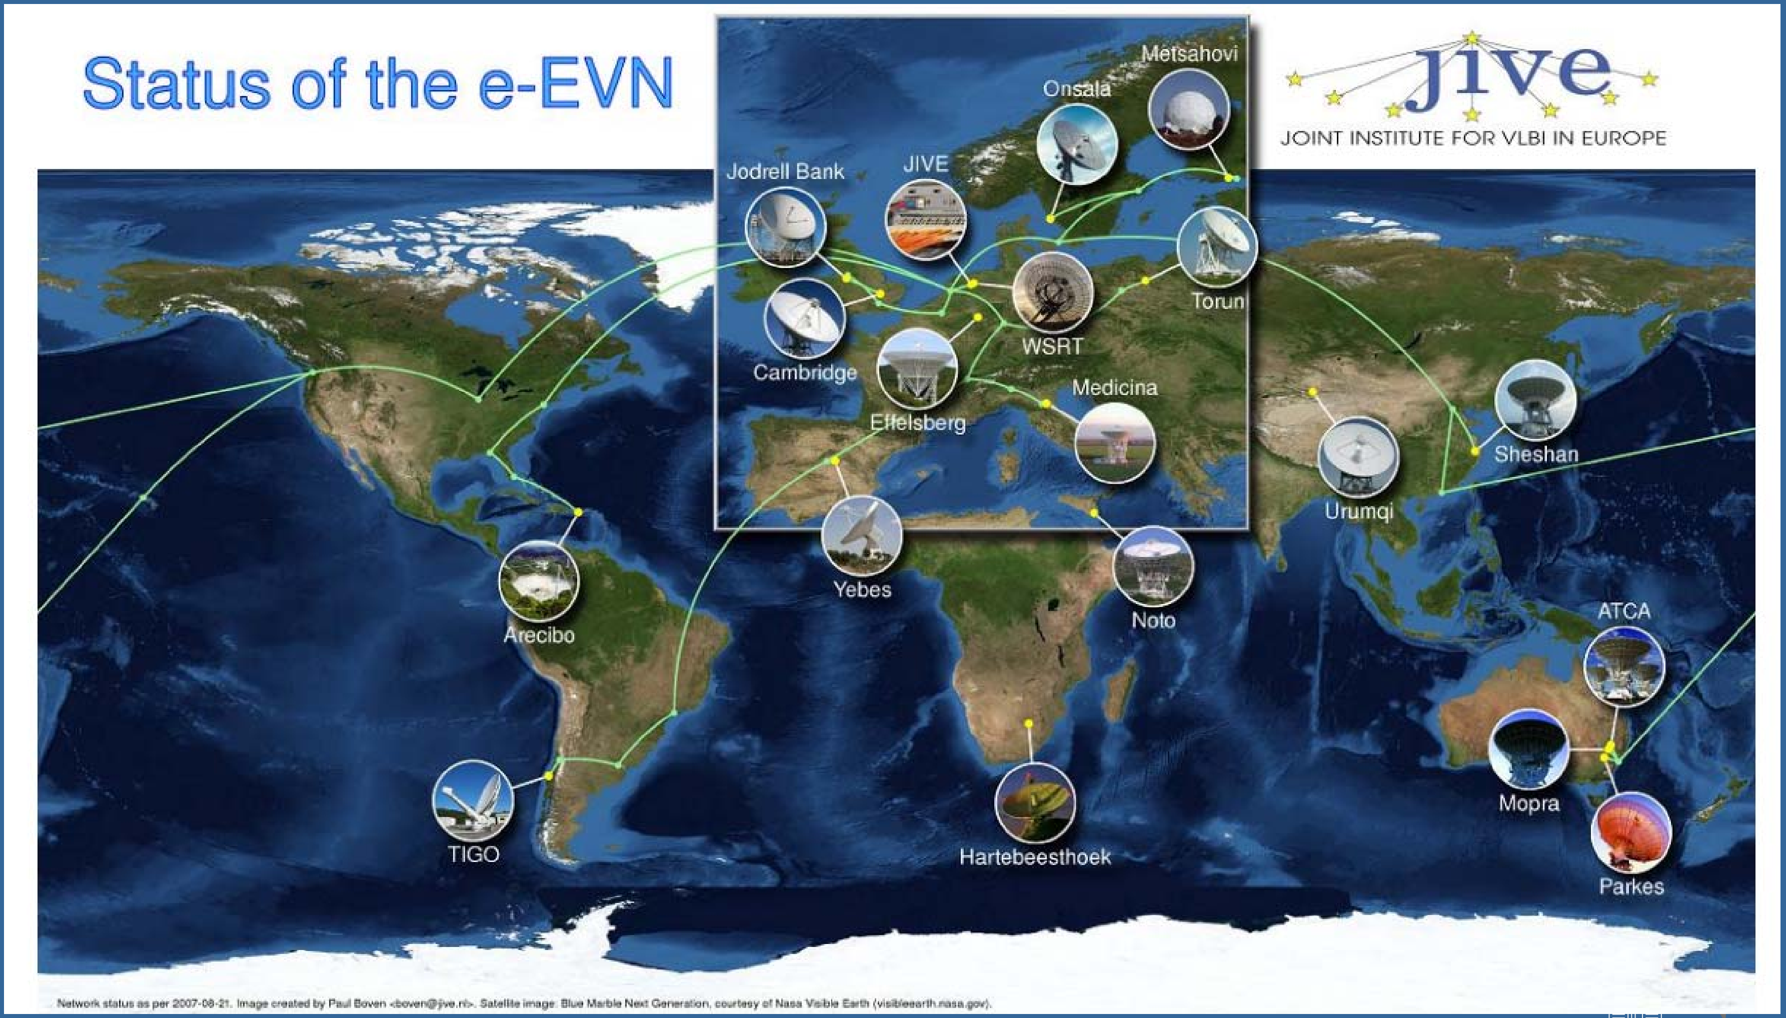

Even connections to Canary Islands!

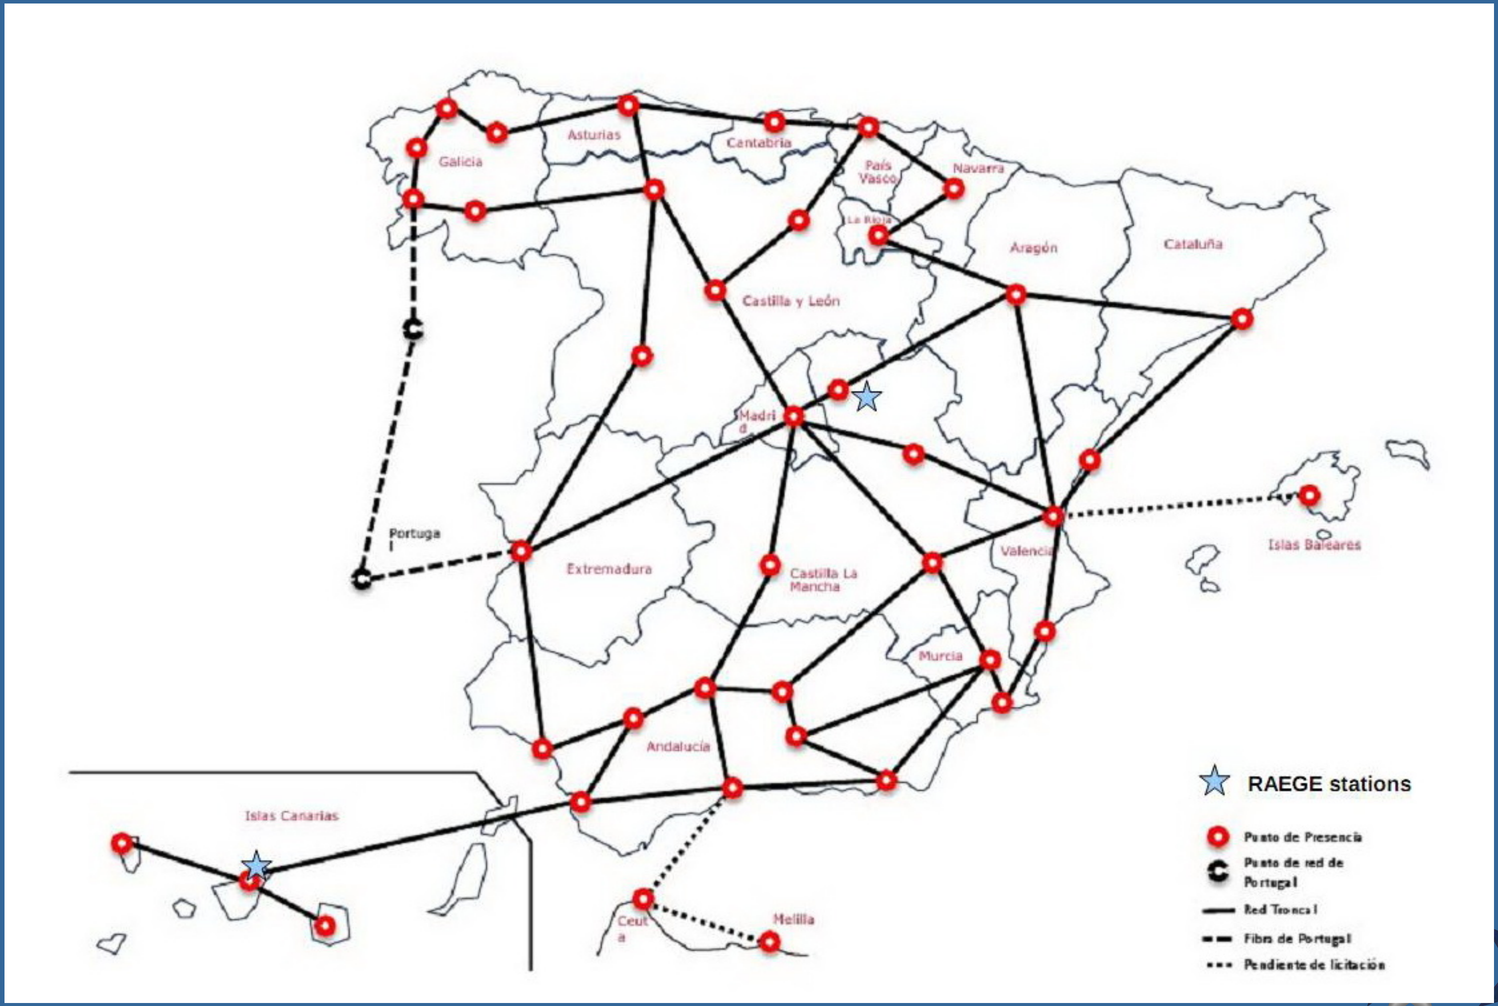

Connectivity rapidly improving to South America

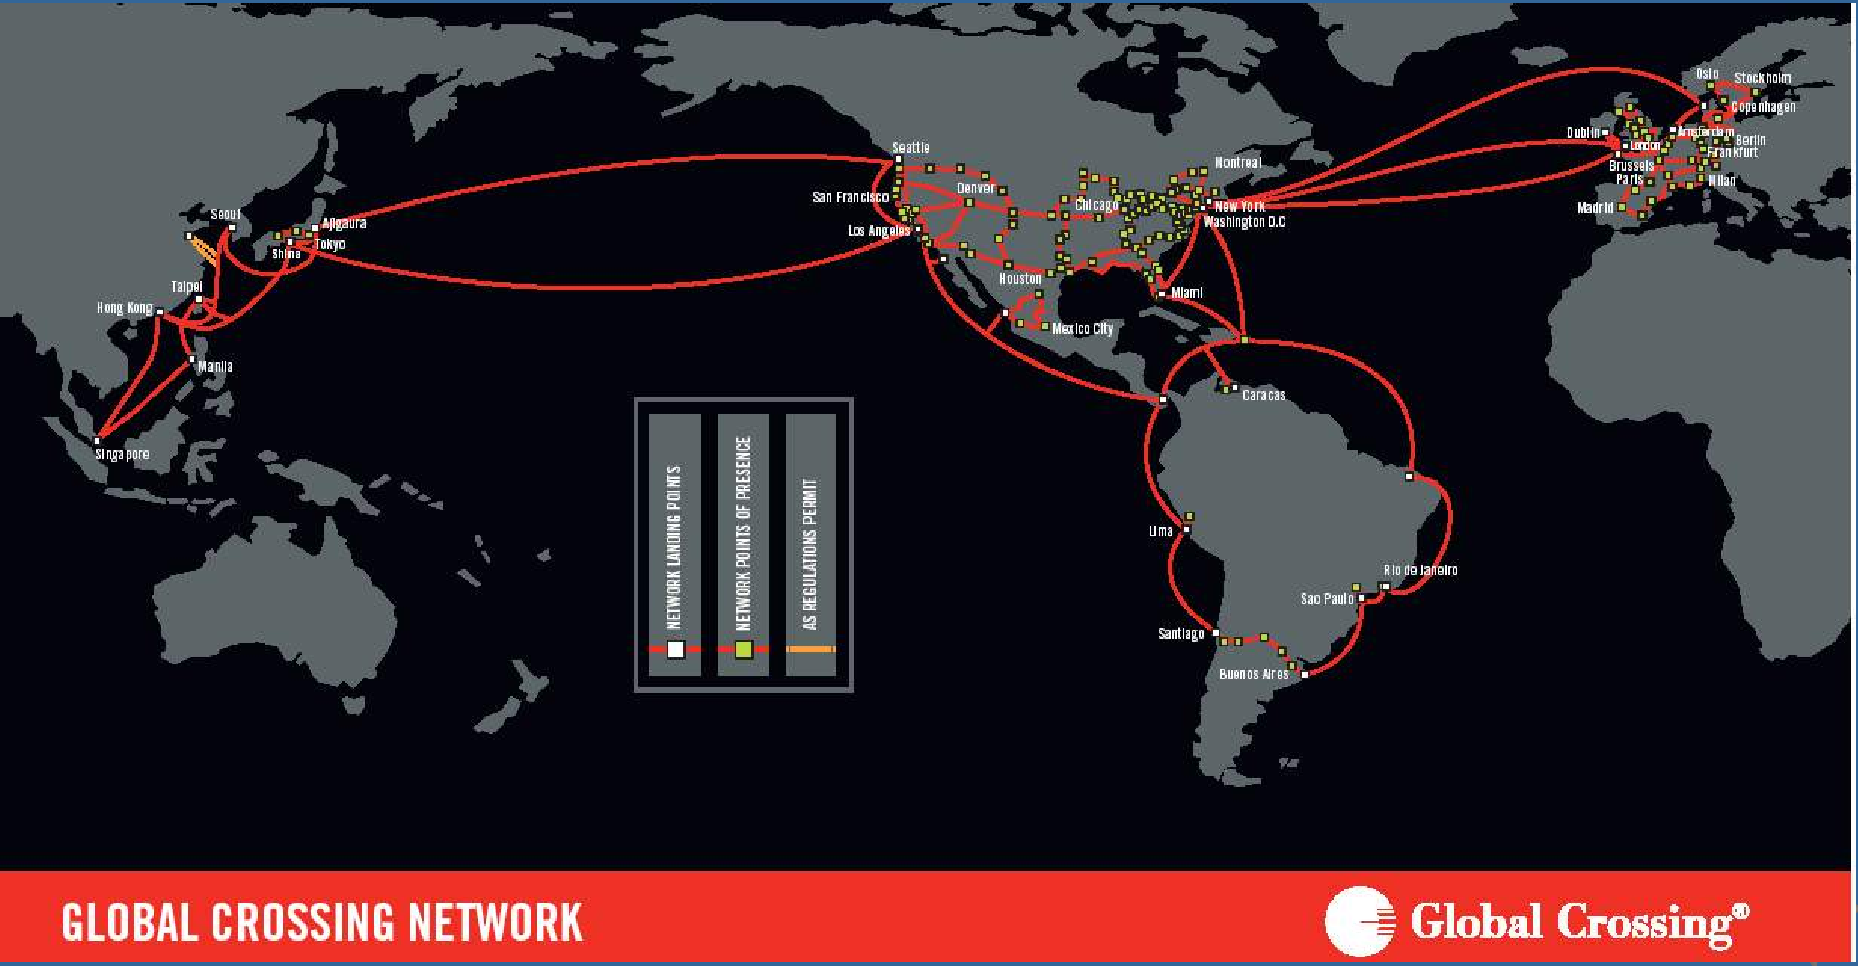

Japan already has many of its telescopes connected at high speeds

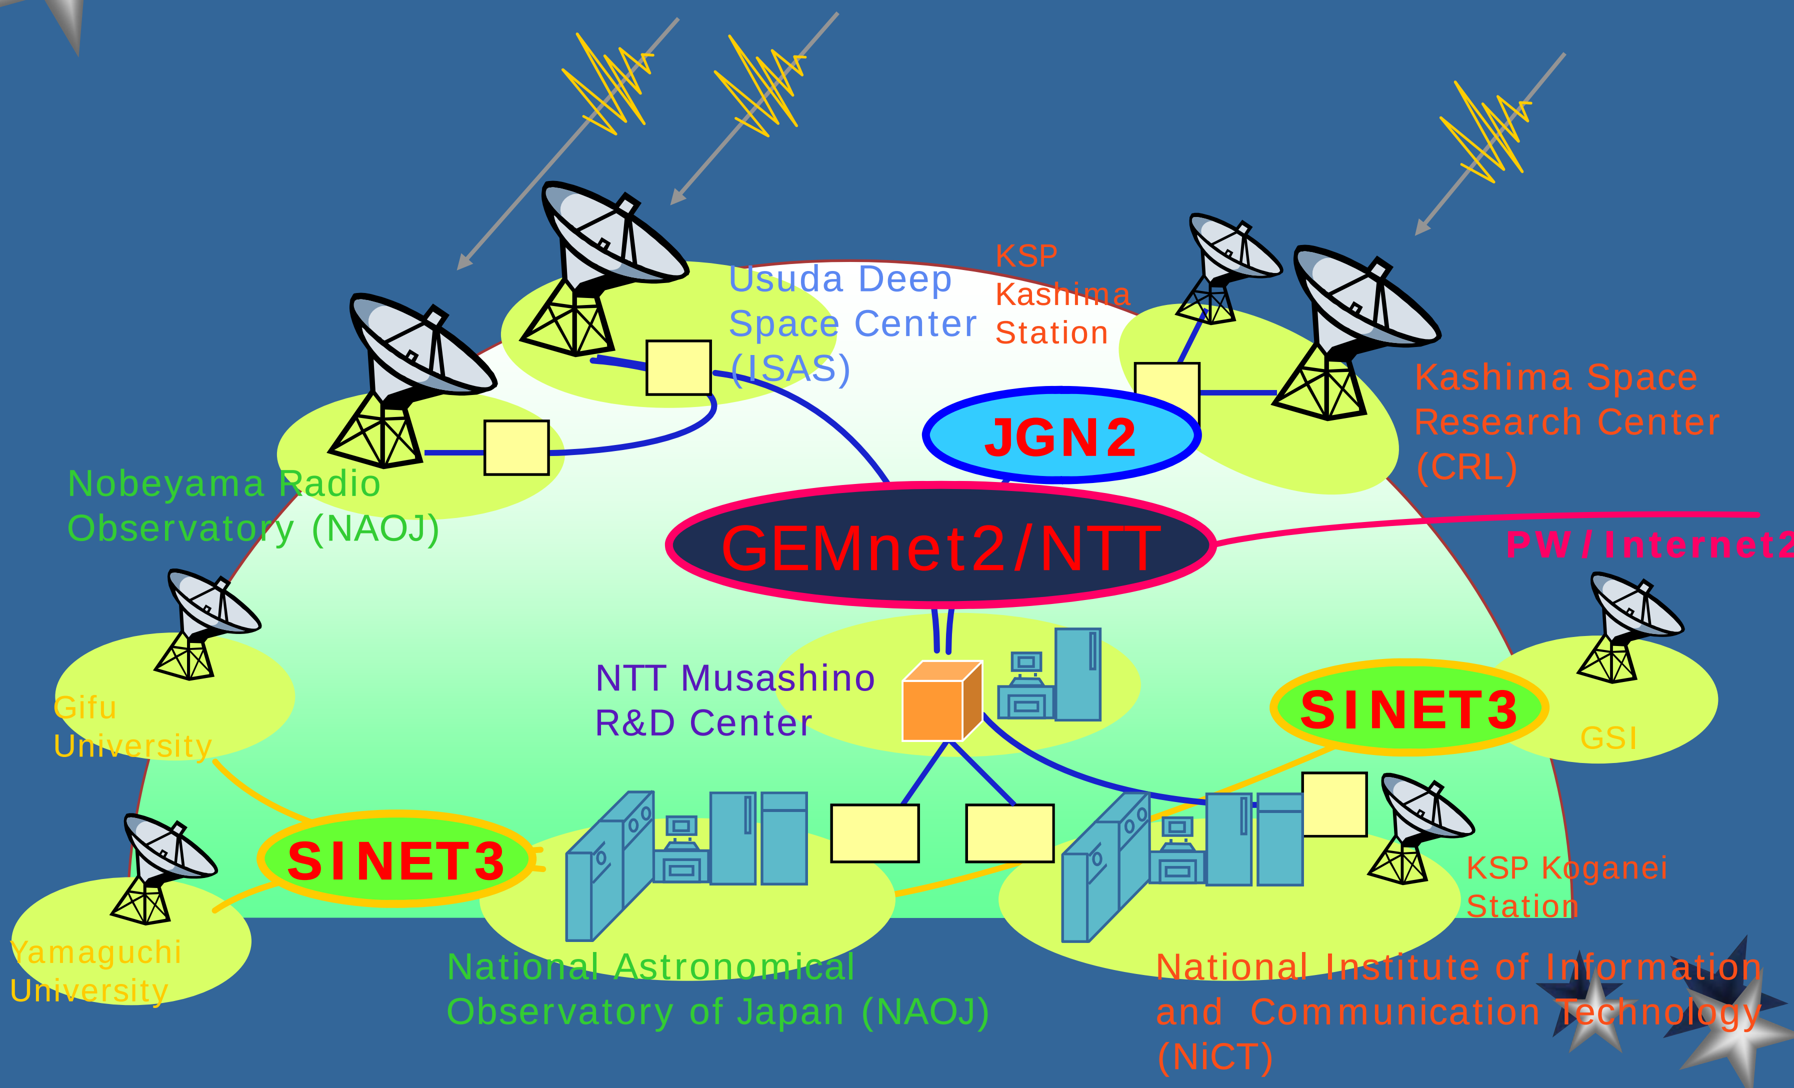

China connections are increasing

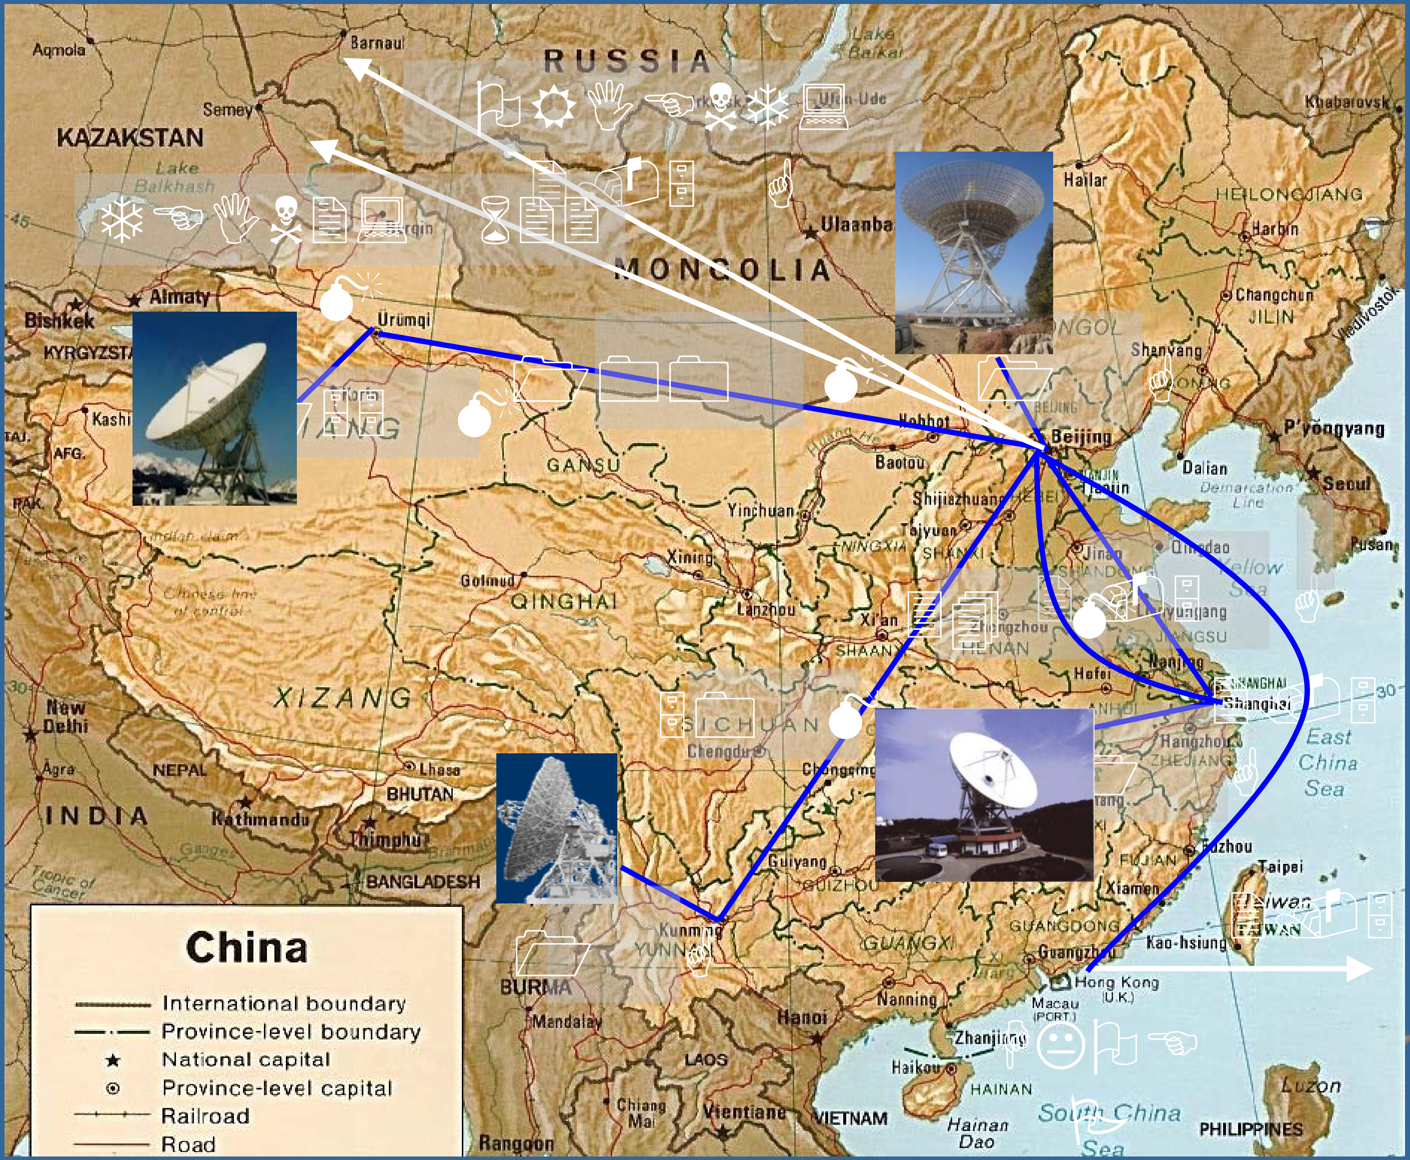

Australia is in process of connecting stations at 10 Gbps

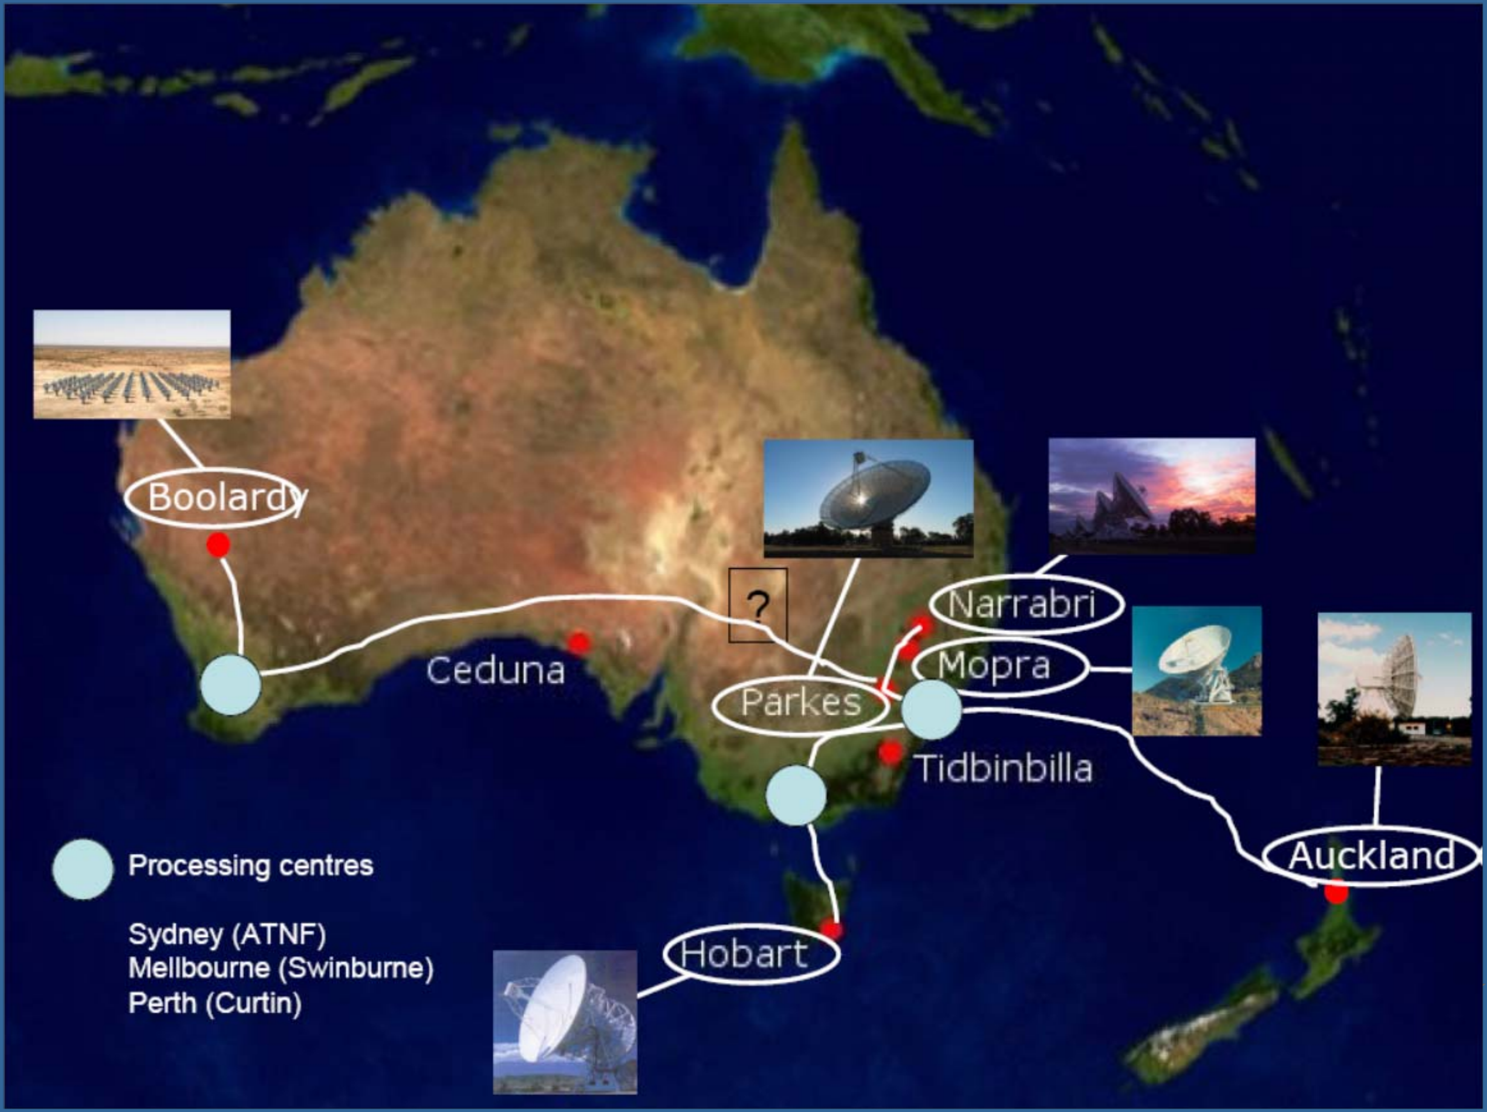

African connections are booming

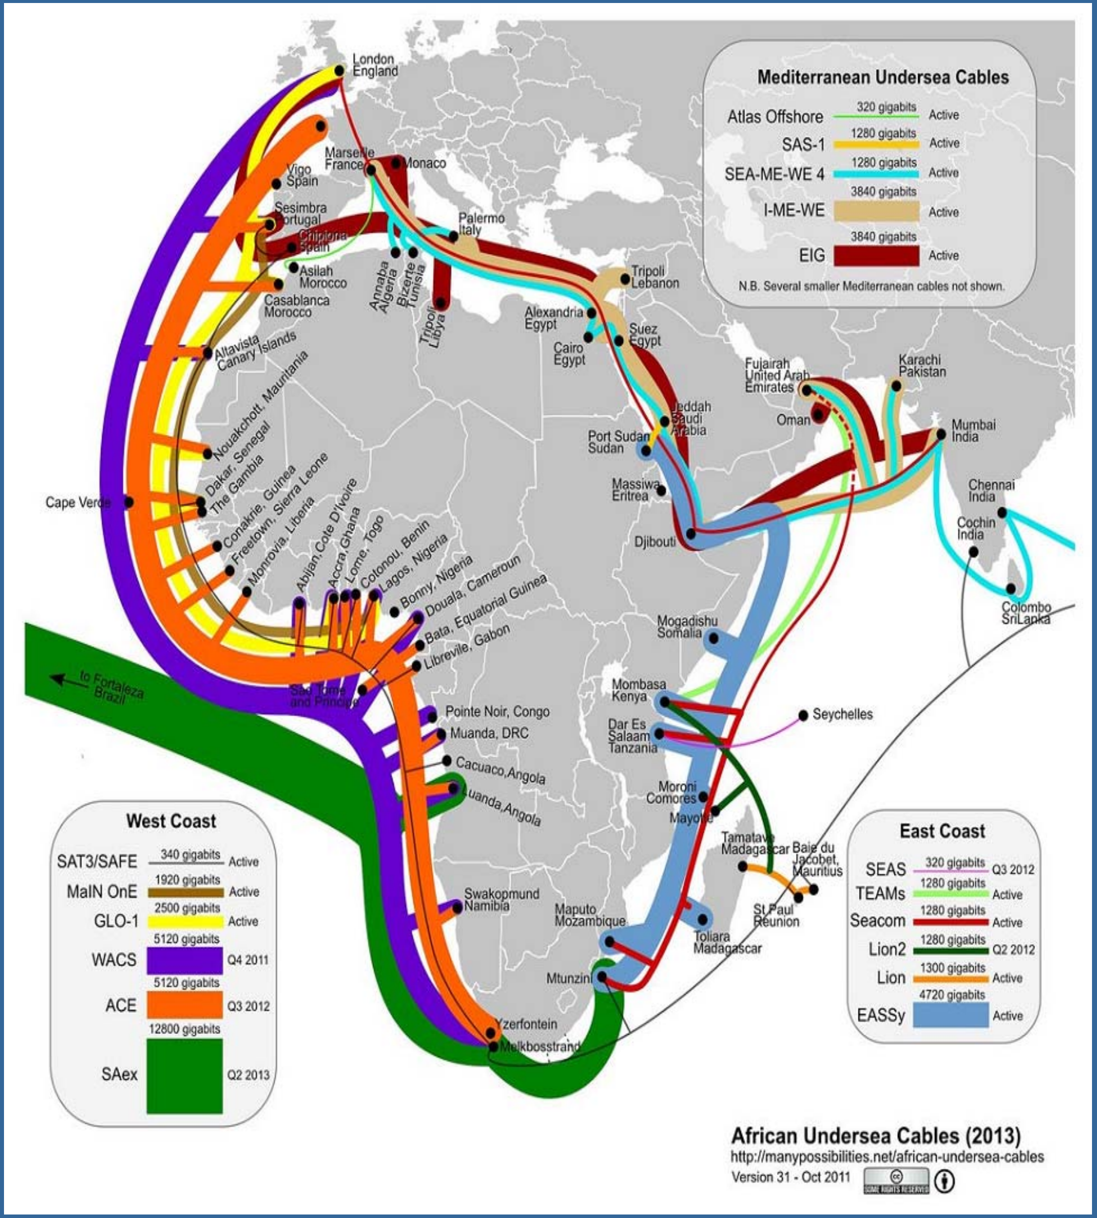

Australian connections

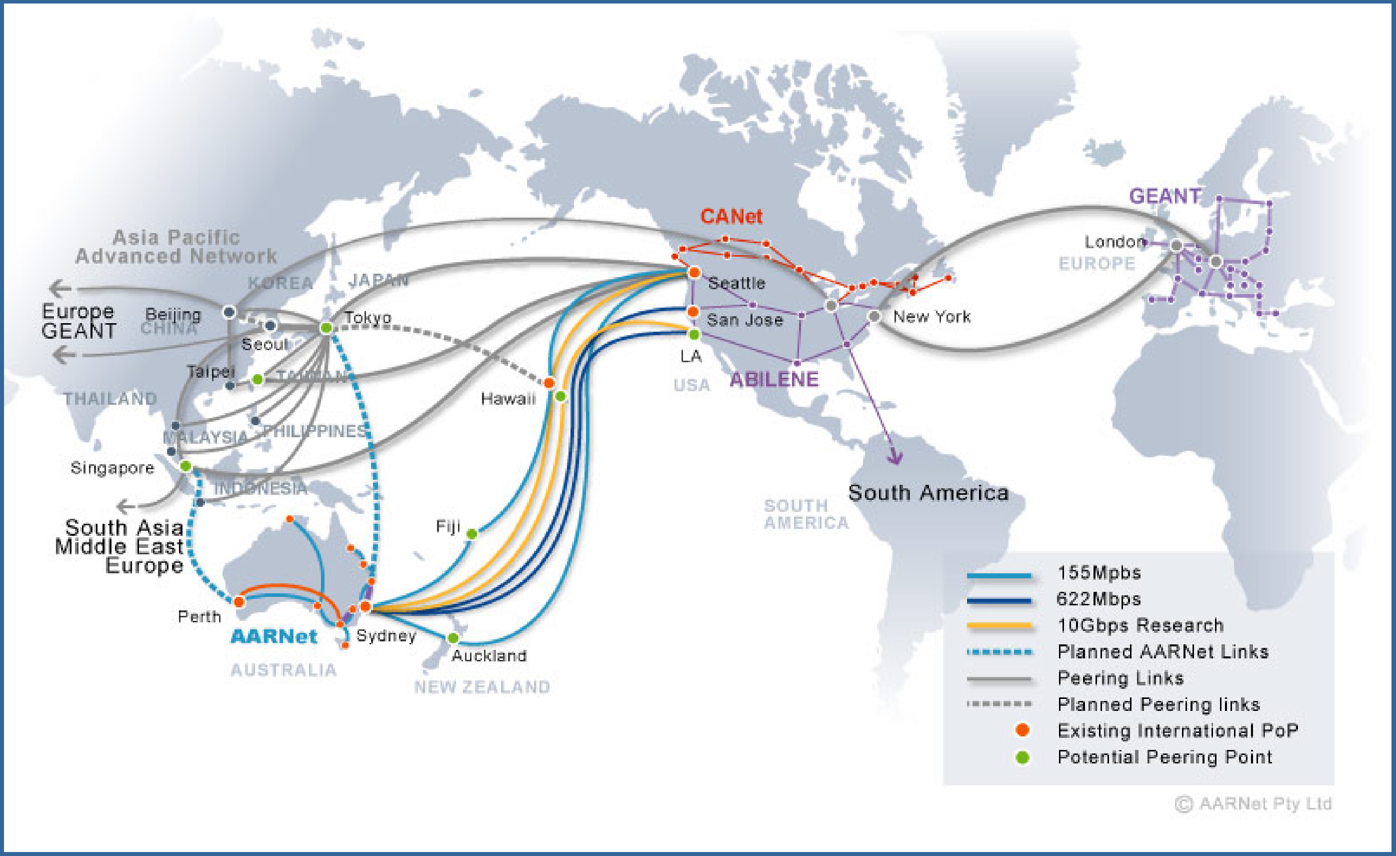

Lots of links across the Atlantic

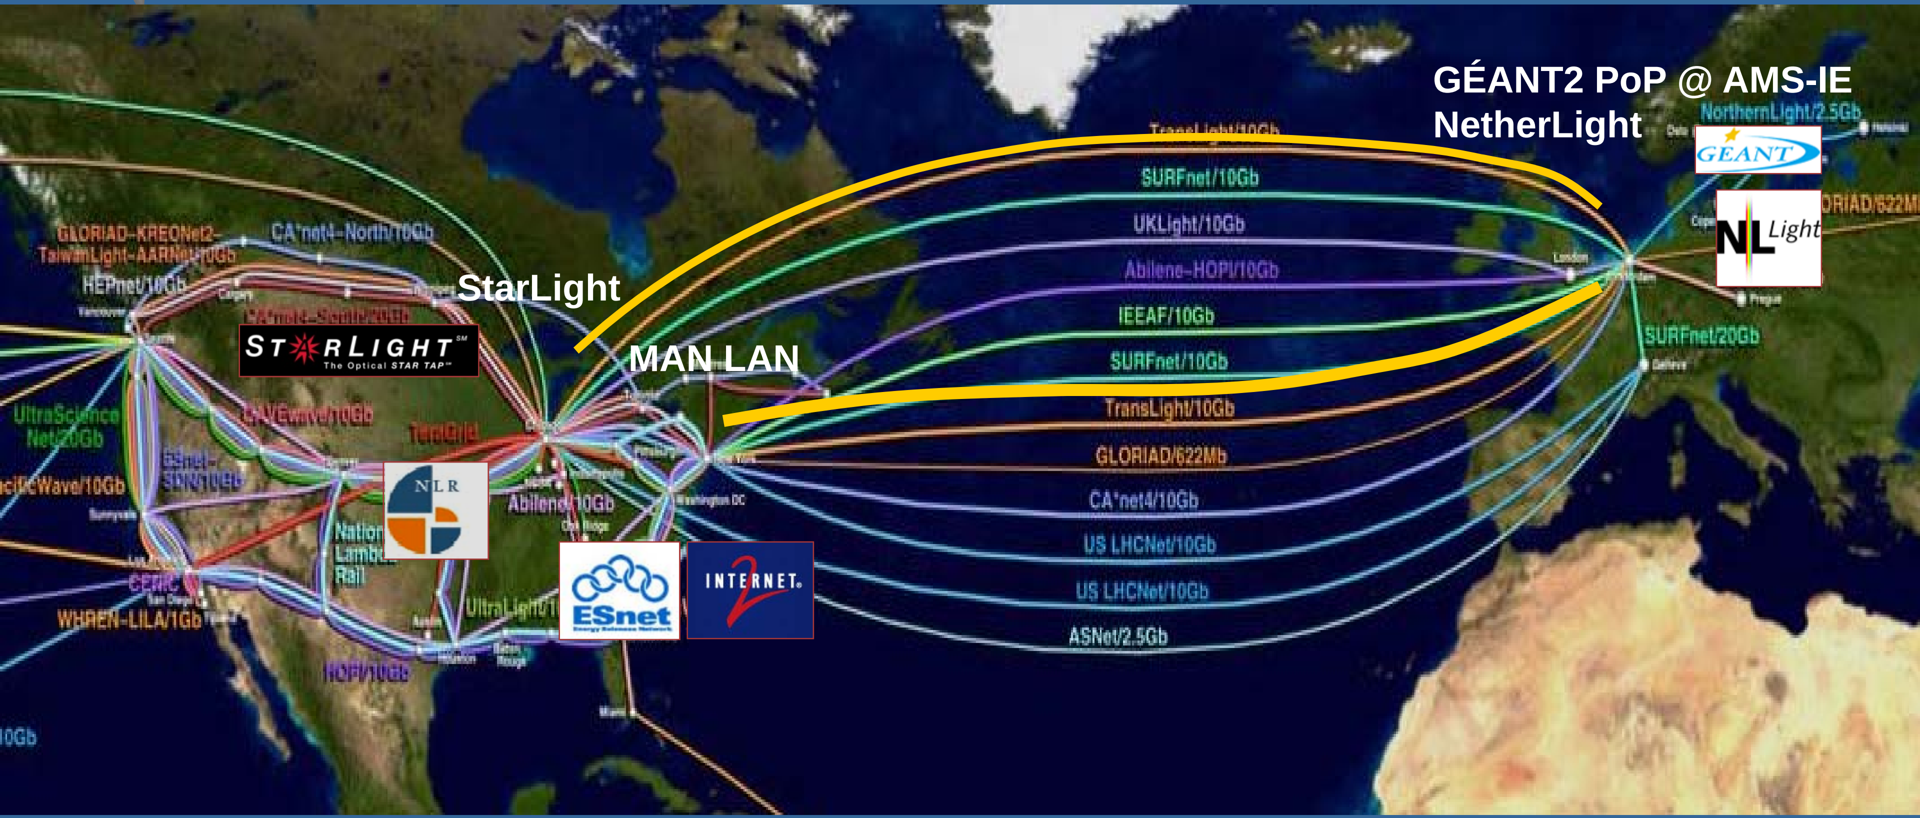

www.startap.net/translight

Lots of links across the Pacific

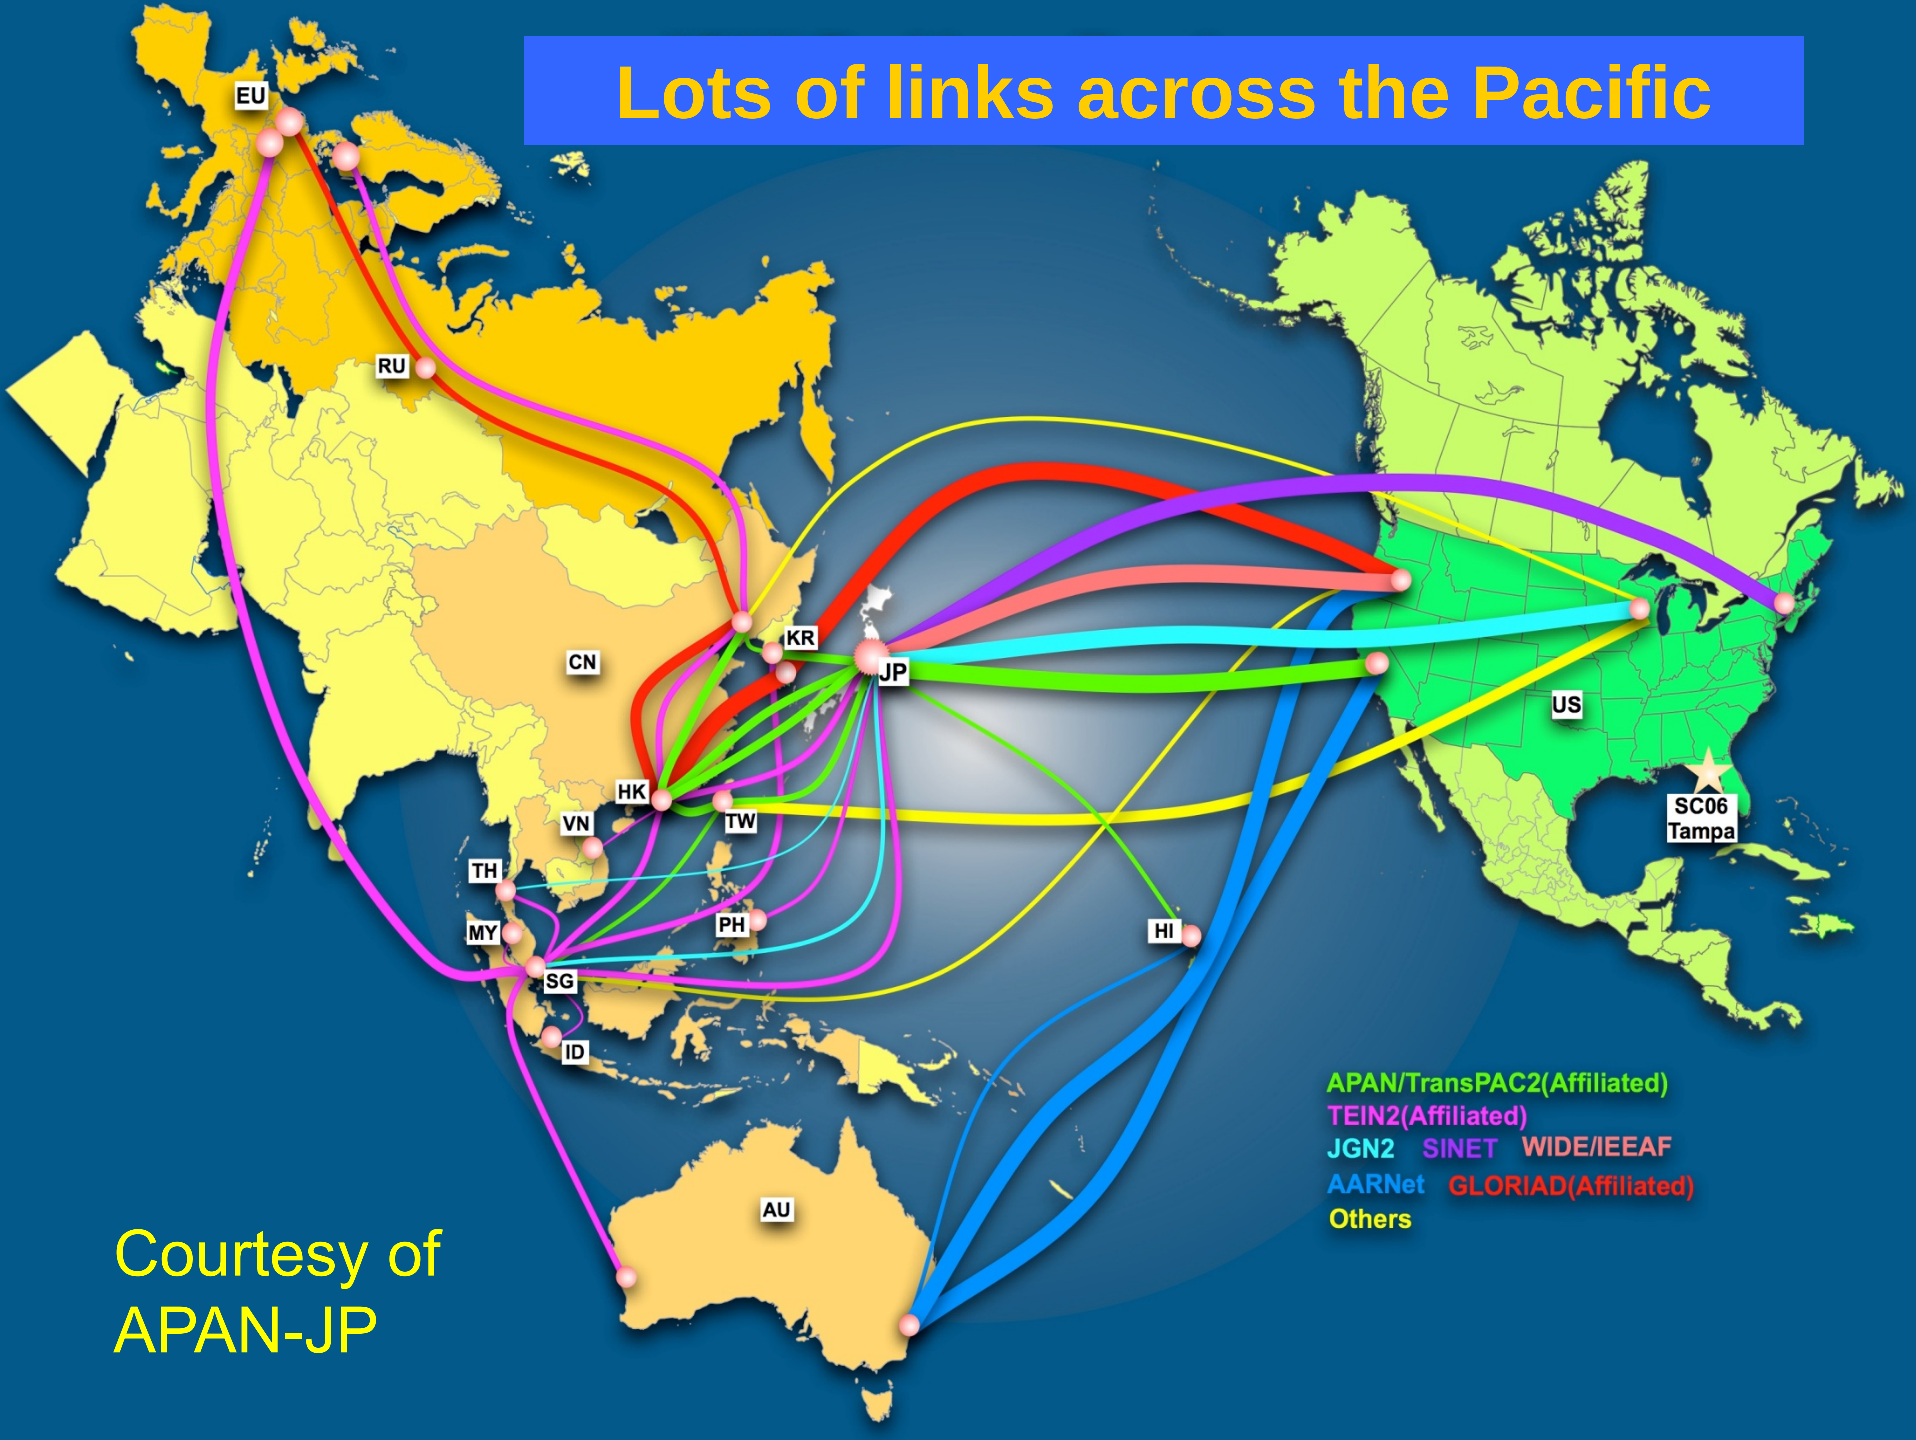

Courtesy of APAN-JP

The 'Last-mile' problem

- Many of the world's radio telescope were deliberately built in remote locations!
- As a result, most of the world's telescopes are not well connected
- Direct fiber cost is relatively low– $60/fiber-km in 80-fiber bundle
- But –
  - Europe: >$20/m (or any populous wide-area)
  - U.S.: >$10m (in simplest desert environment)
- The upside: there is developing a lot of momentum and support from the greater networking community to get the job done!

VTP
(VLBI Transmission Protocol)

- Companion specification to VDIF
- Specifies e-VLBI data-transmission protocol for VDIF-formatted data streams
  - Normally must use UDP or UDP-like protocol to maintain necessary data rate
  - Addition of Packet Serial Numbers (PSNs) helps to keep packets organized and identify missing packets
    (a few missing packets are no normally a problem

e-VLBI Network Monitoring

- R&E networks are always works in progress
- Unfortunately, can’t always rely on network readiness
- Automated monitoring allows view of network throughput variation over time
  - Highlights route changes, network outages
  - Helps to highlight any throughput issues at end points:
    - E.g. Network Inteface Card failures, Untuned TCP Stacks

e-VLBI Conclusions

- e-VLBI is critical to the success of VLBI2010
- e-VLBI is riding an unprecedented wave of global network connectivity and networking community enthusiasm
- 10-100 Gbps/antenna is technically possible with e-VLBI
- Near-real-time data transfer and record is a potential interim solution for less than full time observations
- But….there are e-VLBI challenges!
  - Last-mile connectivity and network costs remain a challenge for many

e-VLBI challenge –
a B747 loaded with recorded digital media!

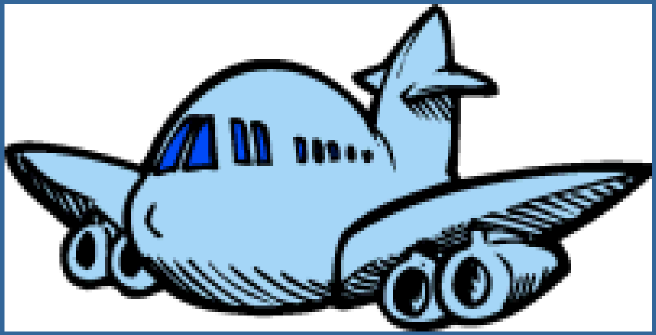

Payload: 140 tons ≈ 140,000 1-TB disks = 140 PB
Based on 24-hr flight time, bandwidth is ~10 Tb/sec!
This is 1000x faster than a 10 Gbps link!

In 1970, with 12” open-reel computer tape at 800 bpi,
a B747 could carry only 1.5 TB; bandwidth ~140 Mbps!
This is 3000x faster than a 56 kbps link available at the time.

The Big Challenge:
When can e-VLBI data transfer capacity catch up to a B747?!

Thanks for your attention

Questions?