



Realizing Einstein's Universe

Unlocking the potential of Radio Astronomy research with High Performance Networking



Charles Smith

Cisco Research Center

chas@cisco.com

The Quest for sub-millisecond Pulsars

Probing the equation of state of atomic matter


- Pulsars are the collapsed cores of once-massive stars some 500,000 times heavier than the Earth but only 20 km in diameter.
- One cubic centimeter of their material weighs more than 1 Billion tons!
- At present astronomers have detected these sources rotating at up to 700 times per second.
- New digital instrumentation coupled with a high performance network and massive supercomputers would open vast regions of parameter space within the observations that may contain proof of sub-millisecond pulsars.
- The discovery of a single sub-millisecond pulsar would result in the elimination of over half the current viable equations of state of nuclear matter.

millisecond pulsar



Finding that Gravity Wave

The last “Holy Grail” of Physics

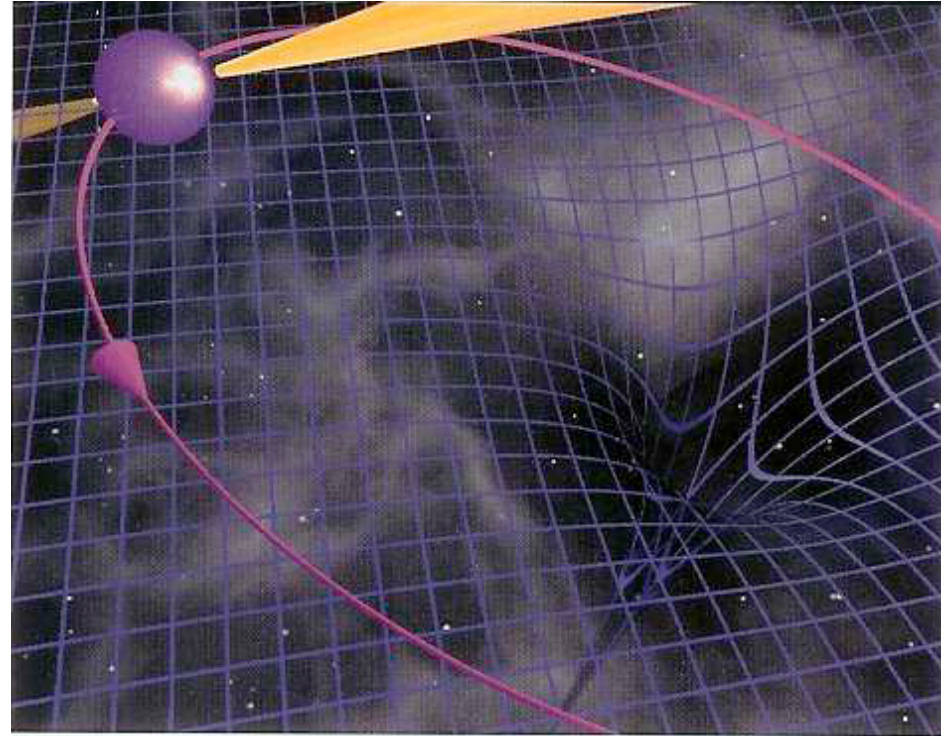
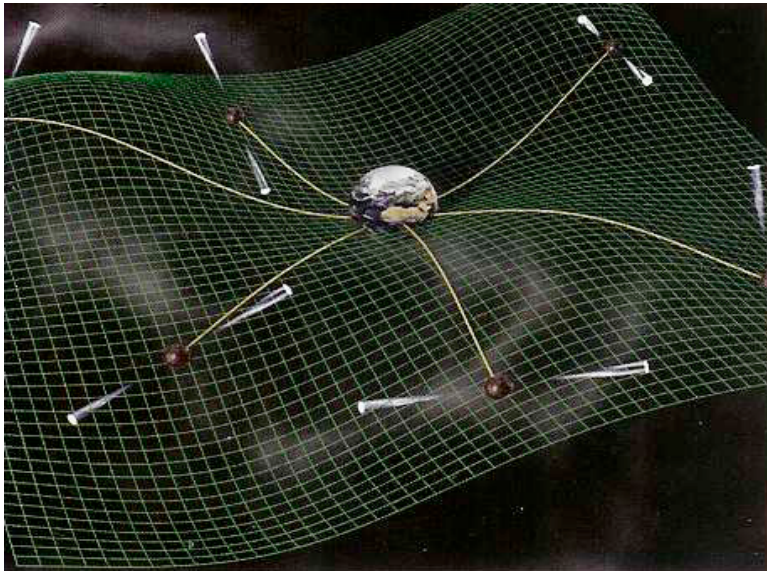


- Super-massive black holes permeate the Universe with a gravitational wave background that physicists seek to detect with a religious zeal.
- Detection of these waves is possible by measuring space/time distortion of distant pulsing stars with giant radio telescopes.
- Linking supercomputing resources at giant radio telescopes eliminates instrumental artifacts that have hitherto prevented astronomers from achieving the theoretical limits imposed by Nature.
- High performance networks now have the capability to deliver the requisite bandwidth required for the most sensitive gravitational wave detector ever built.

Pulsars - The “Universal Clocks”

- Pulsars are the Universe’s best clocks.
- Using the timing of Pulsars Gravity wave’s and precise proofs of extreme cases General Relativity can be attained.
- e.g. Pulsar 1909 Period

$P = 2.947108021647488(6) \text{ ms}$



Finding ultra-high energy neutrinos


New realms in particle physics

- Man-made particle accelerators are dwarfed in power by exploding stars which liberate $>10^{46}$ Joules in milliseconds.
- This gives rise to single elementary particles that have more energy than 10^{20} electron volts (ie more than a Joule / particle!).
- Interaction with the cosmic microwave background produces ultra-high energy neutrinos that are yet to be detected.
- GigaHertz data streams linked to supercomputers provide a new mechanism for detecting these postulated particles as they interact with the Moon and radiate Cerenkov radiation in the radio band.



Crab Supernova Remnant

*Parkes 70m
And the Moon*




The Transient Radio Sky

A New Class of Astronomical Source

- Large-scale surveys of the sky have recently revealed a new class of impulsive transient radio sources.
- Limited frequency and time resolution by analogue spectrometers leave this new area of science in its infancy.
- Development of new digital systems with 10 Gb data streams open new realms of parameter space impossible to penetrate without utilizing massive computational resources networked directly to the most powerful telescopes.



*2 microsecond burst
of radio emission discovered
on the Swinburne Supercomputer
from a millisecond pulsar
6000 light years distant*



Realizing Einstein's Universe


Current Technologies & the move to Software Correlation

CPSR2 Parkes Australia

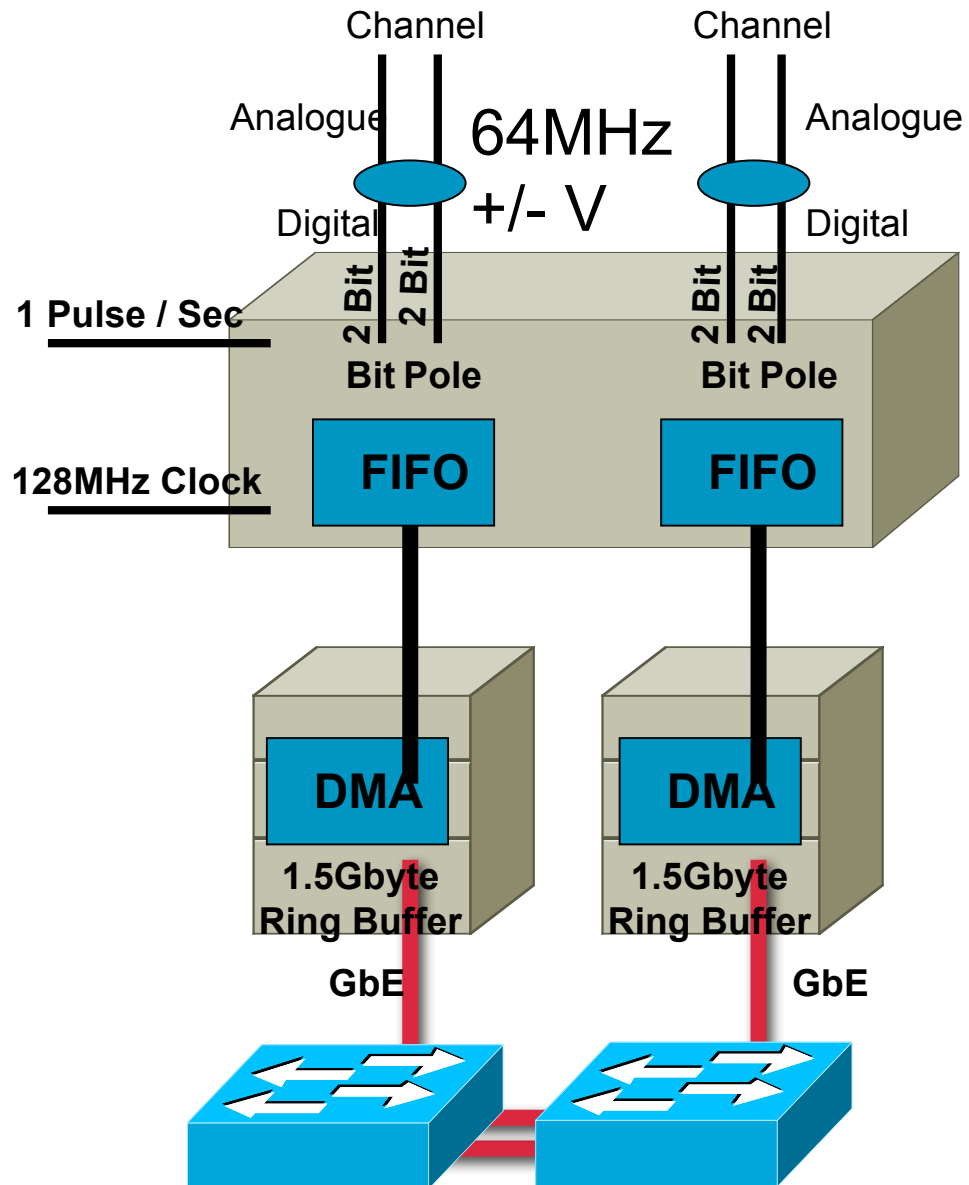


| | |
|-------------|-------|
| Nynquist | 2 |
| Channels | 2 |
| Sample Rate | 64MHz |
| Bits | 2 |
| Polarities | 2 |


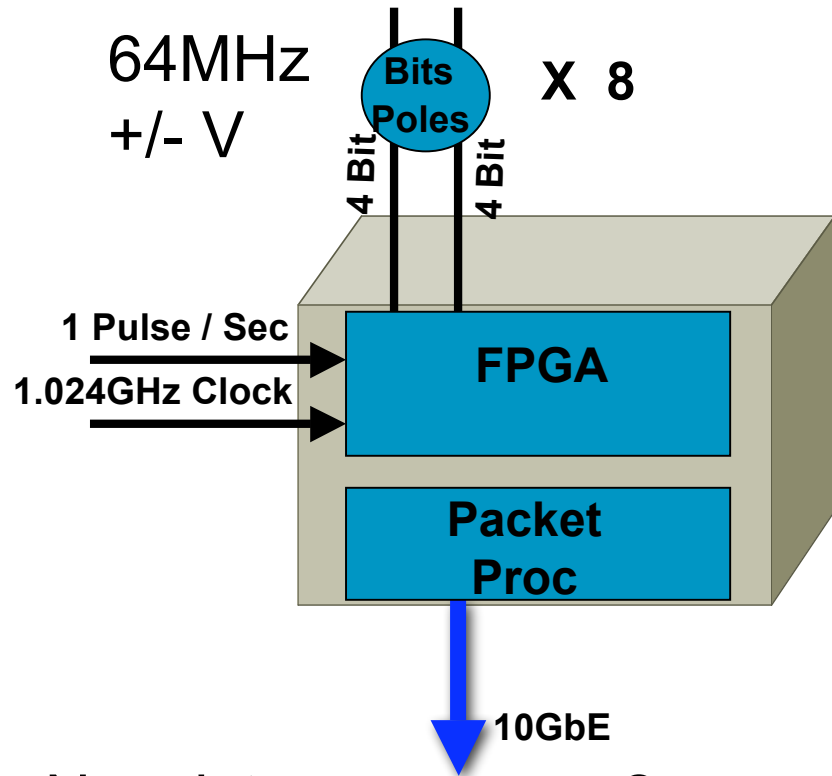
= 1024 Mbit ~ 2GbE



CPSR2 Parkes



Nextgen Data Acquisition



Nyquist Channels 2
 Channels 8
 Sample Rate 64MHz

Bits 4
 Polarities 2

8192 Mbit ~ 1 x 10GbE or 16 x GbE

Realizing Einstein's Universe

The Square Kilometer Array

The Square Kilometer Array - SKA

Parameters

Design Goals

| | |
|---|---|
| Frequency range | 100 MHz to 25 GHz |
| Sensitivity (area / system temperature) | 20,000 m ² /K (or G/T - 68 dB/K) between 0.5 and 5 GHz, 10,000 m ² /K up to 25 GHz |
| Field-of-view | 1 square degree at 1.4 GHz and 200 square degree at 0.7 GHz; target is 4 simultaneous FOVs, each with full sensitivity |
| Angular resolution | <0.1 arcsecond at 1.4 GHz |
| Instantaneous bandwidth | 25% of band centre, max 4 GHz |
| Spectral (frequency) channels | 10,000 per band per baseline |
| Calibrated polarization purity | 10,000:1 |
| Synthesized image dynamic range | >1,000,000:1 at 1.4 GHz |
| Imaging processor computation rate | 10 ¹⁵ operations per second |
| Telescope output data rate | 1 terabyte per minute (typical) |

TIMELINE

Concept

1994

International Working Group

1995

Start of Prototyping

2000

Signing of first Memorandum of Agreement

2005

Signing of extended Memorandum of Agreement

2006

Site Ranking Decision

2009

Final Technology Decision

2010

Construction of Phase 1 on site

2013

Early Science

2014

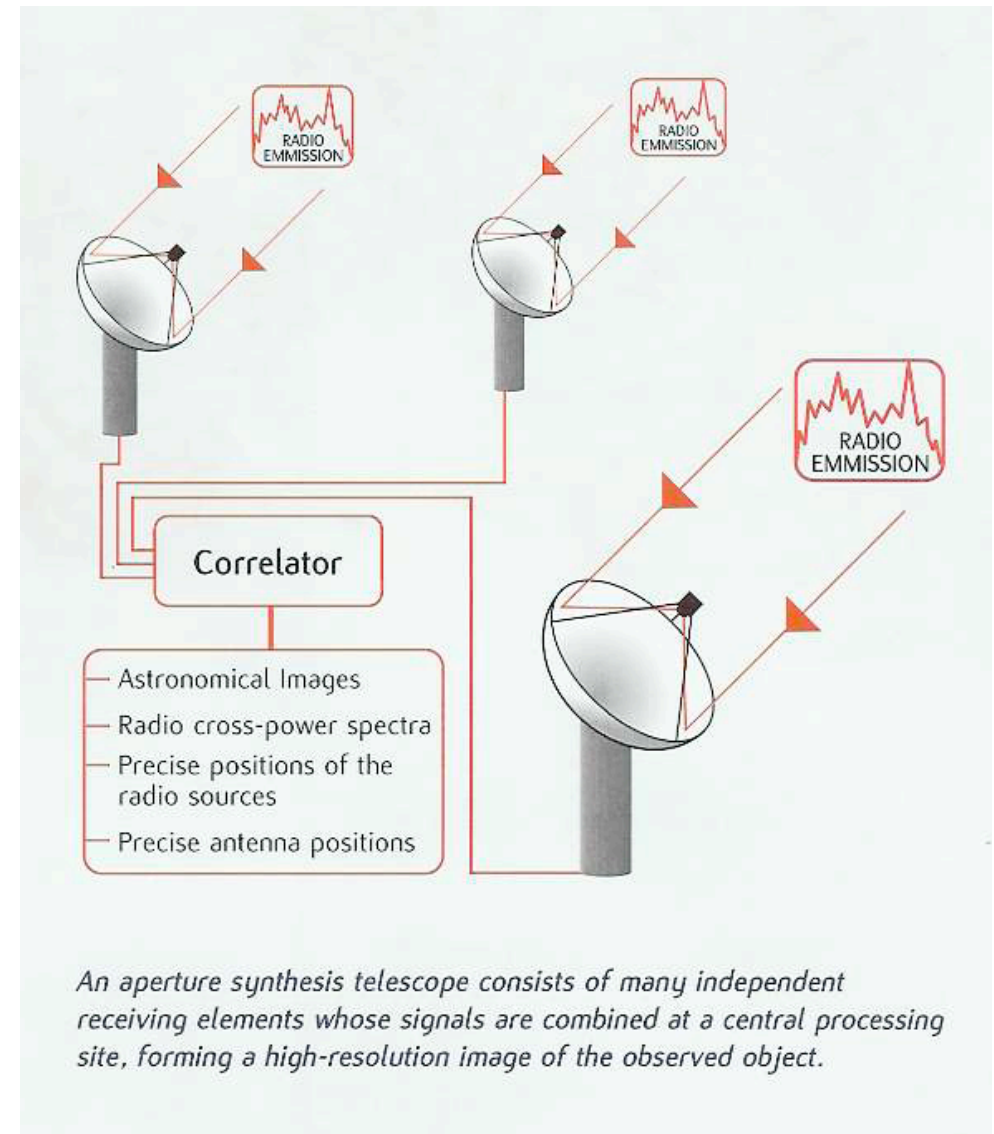
Construction of full array

2020

Full SKA operational!

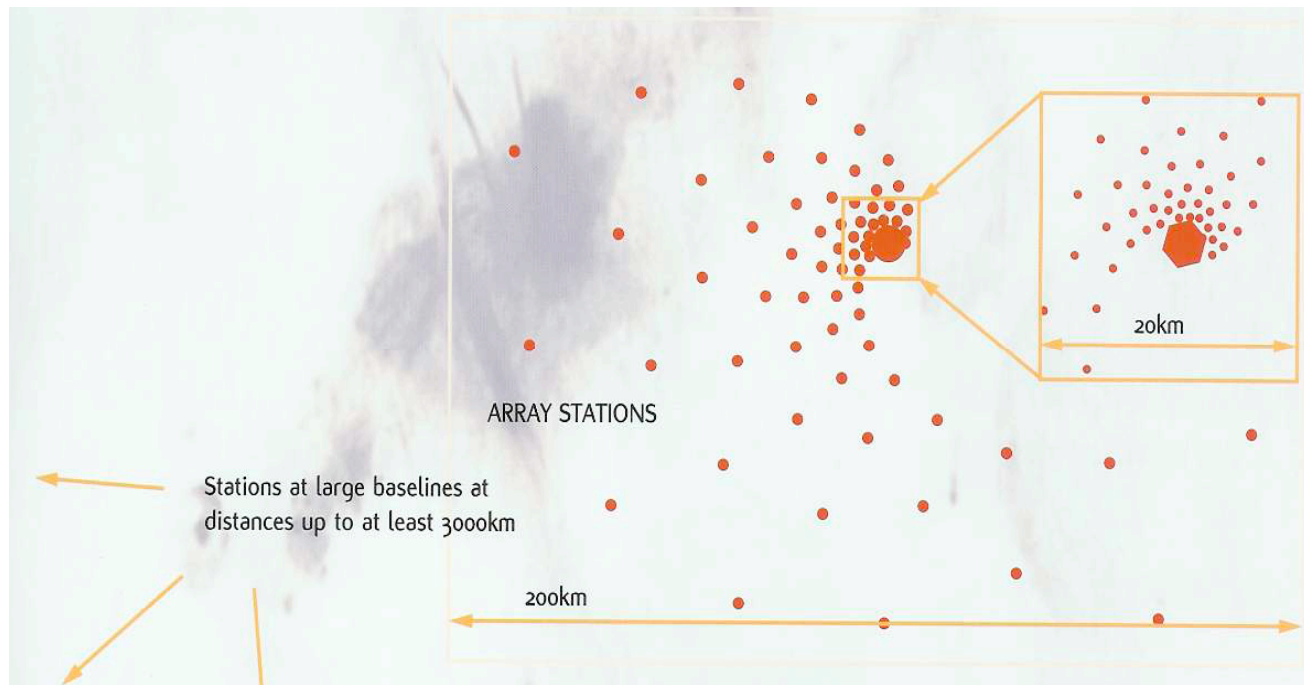

Aperture Synthesis Arrays

- Aperture Synthesis telescopes consist of many receiving elements whose signals are combined at a central processing site.
- The combination and correlation of these signals is a relatively simple but highly compute intensive Fast Fourier Transform ideally suited to parallel computing environments.



SKA Central Site Architecture

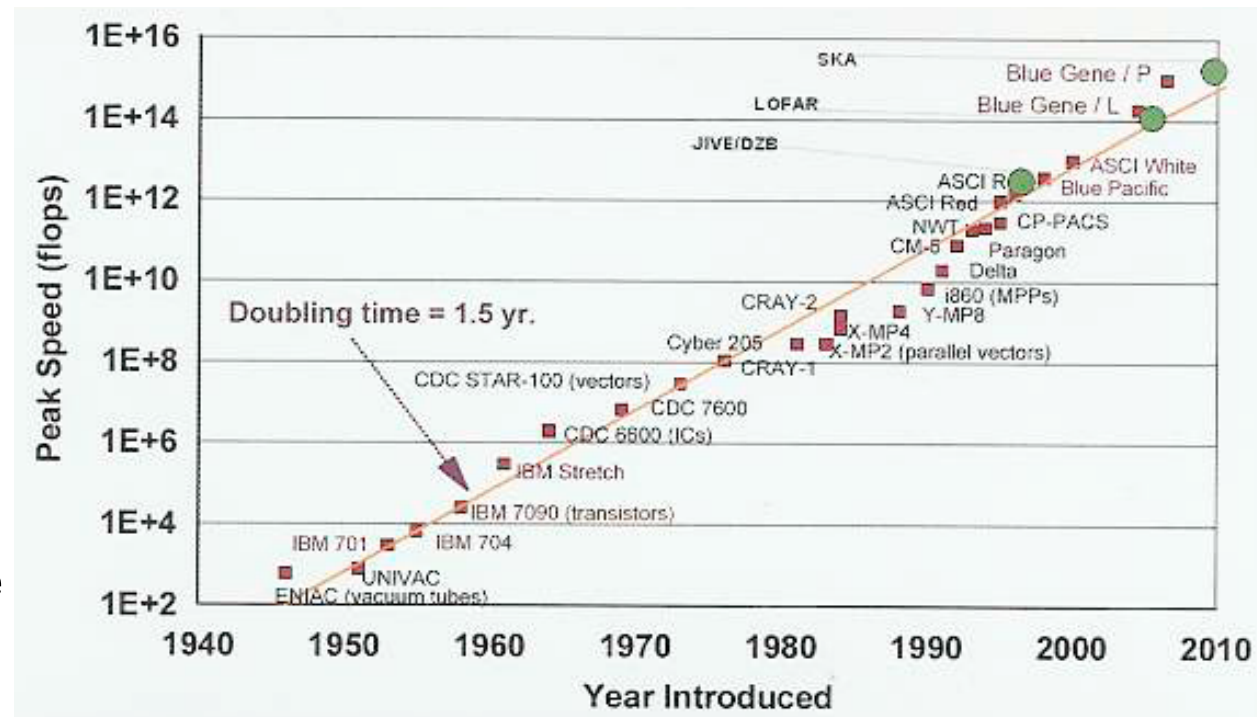
- Possible distribution of receiving stations on small and moderate distances from the core
- Further stations will be located at distances up to at least 3000km in a logarithmic arc topology from the core.



Engineering Challenges


- Low-cost collecting area
- Low-noise, highly integrated receivers
- Phased-array antenna technology
- Wideband optical-fibre signal transport
- Fast, high resolution, A-D Converters
- High performance computing environments
- Data storage & innovative retrieval technologies

- Can Moore's law keep up ?
- It's a lot more than simply how fast your CPU is...
- This is continental scale real-time, distributed compute environment

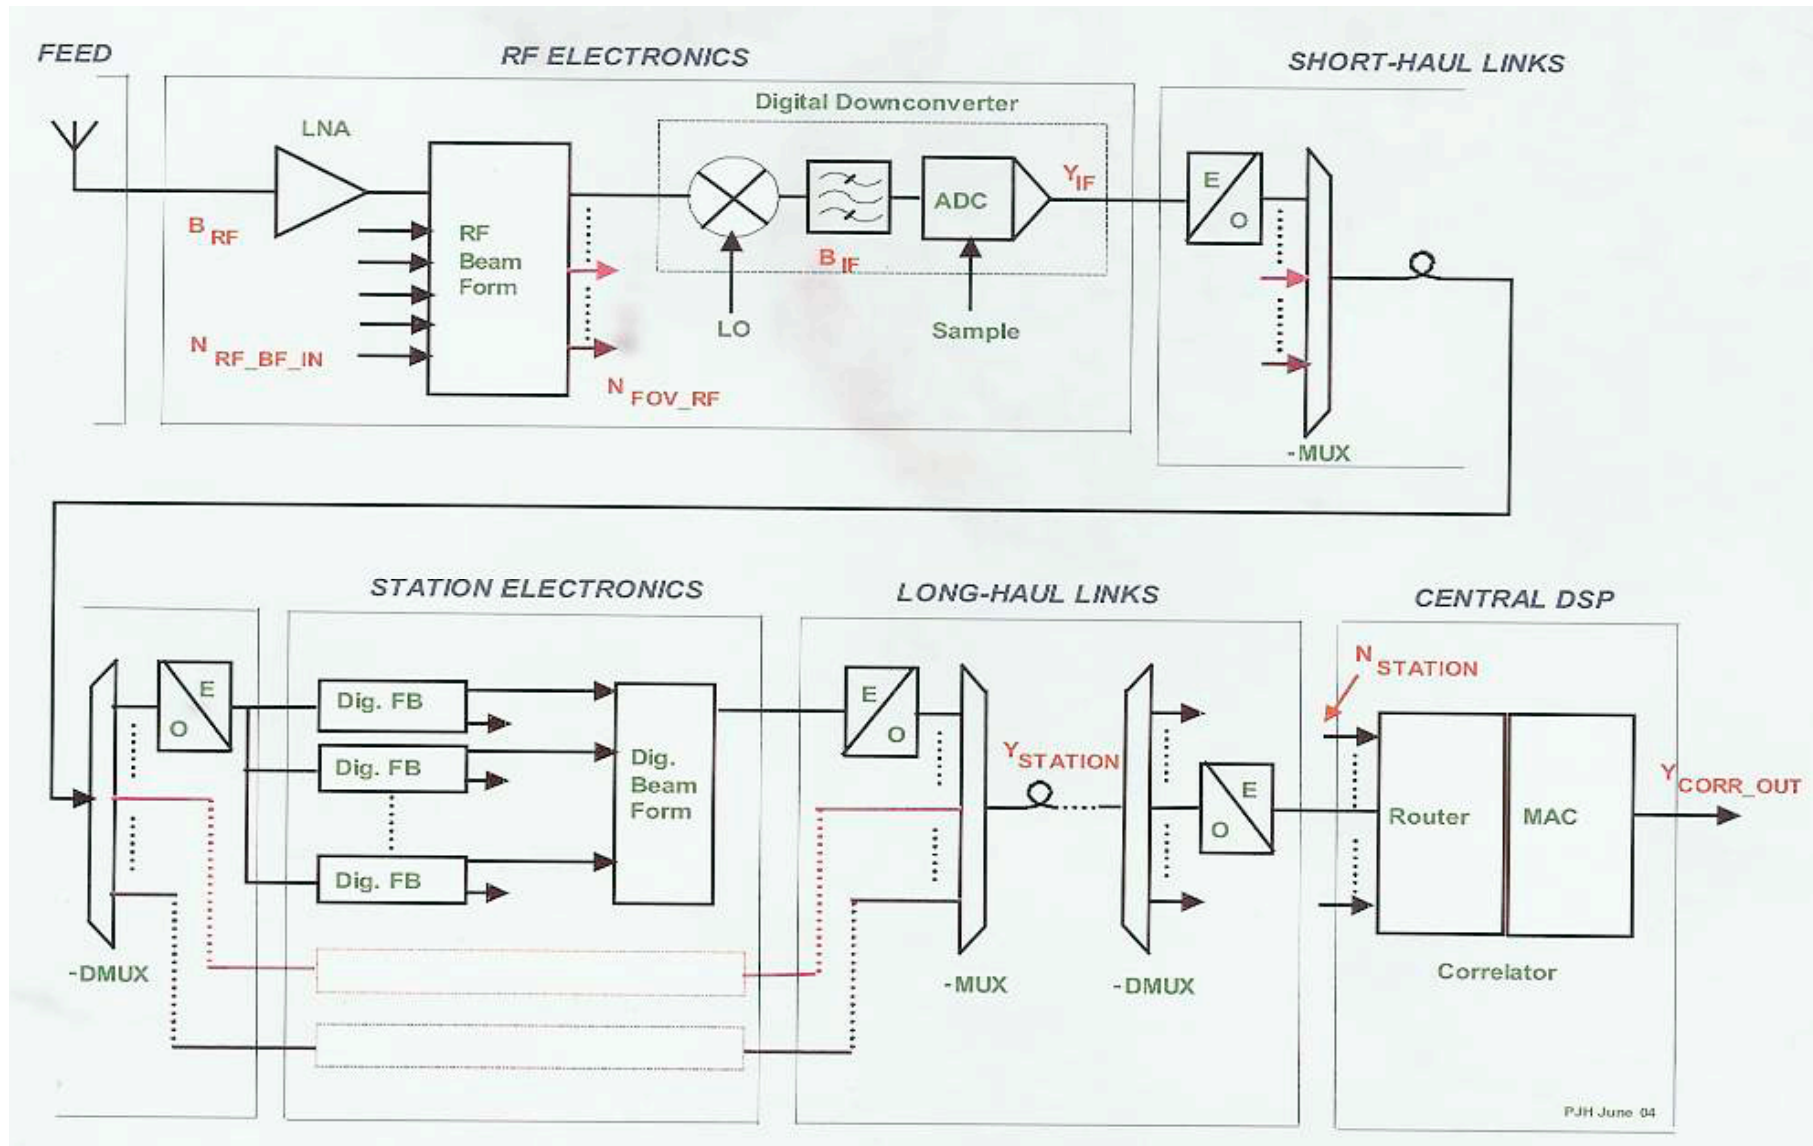


Antenna Design Considerations

- Aperture arrays for observing much of the whole sky at one. Great for surveys.
- High frequency dishes 12-15m (Allen telescope uses these)
- Small dishes implemented in a phased-array at their foci

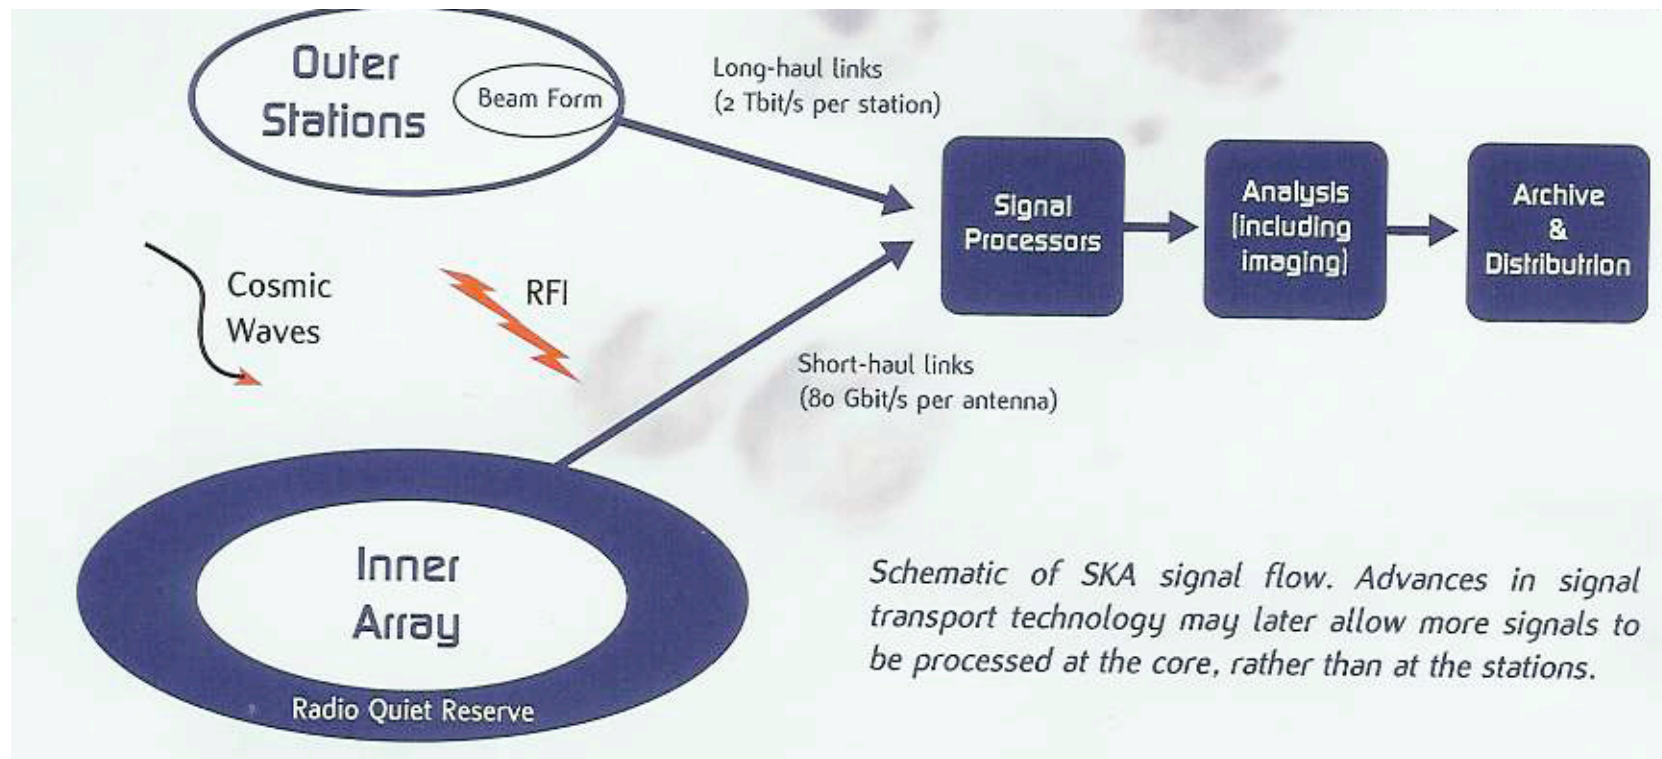


Proposed Architecture



Data Rates

- Undergoing fast, high resolution sampling (10Gb/s) data will be transported to a station processor.
- Short-haul links in the inner array will transfer 80 Gbit/s per antenna
- Long haul links will need a capacity of 2 Tbits / station
- Transfers world-wide of 1 Tbyte images / minute



Realizing Einstein's Universe

Lets Crunch some numbers and have a closer look at this from a
Networking perspective.

Sensor Collection

Each Sensor Gathers:

8 bits of Sampling

2 bits of Polarisation

2 bits of Nynquist Sampling

Samples at the rate of 256 samples / u second

Note this may move to 512 samples / u second

$8 \times 2 * 2 * 256 \times 10E6 = 8.19E9 \text{ bits/s} =$

7.6 Gigabits / second

That is a 10GbE Port required / Sensor

Antenna Collection

Each Antenna has 10 Sensors:

There is a focal array of 10 Sensors

7.6 Gigabits / Second x 10 =

80 Gigabits / Second / Antenna

That is a switch with:

8 10GbE Ports Input

2 40GbE Ports Output

Station Collection

Each Station has 60 antenna:

There are 60 antenna per station

80 Gigabits / Second / Dish x 60 =

4,800 Gigabits / Second / Station

That is a storage requirement of:

600 Gigabytes / second

0.6 Terabytes / second

SKA Collection

The SKA has 130 Stations:

There are 130 Stations

4,800 Gigabits / Second / Station x 130 =

624,000 Gigabits / Second

That is a storage requirement of:

78,000 Gigabytes / second

78 Terabytes / second

A typical experiment runs for 3 days:

$78 \times 60 \times 60 \times 24 \times 3 / 1024 = 20,000$ Petabytes of Storage

Reducing the Data Set Size

A First Level of Correlation at the Dish:

There is a 85% reduction of dataset size

7.6 Gigabits / Second x 10 = 33.5 Terabytes / Hour * 0.15

5 Terabytes Storage / Hour

Reducing the Data Set Size

A Second Level of Correlation at the Station:

There is a 85% reduction of dataset size

$$5 * 0.15 = .75 \text{ Terabytes Storage / Hour}$$

That is 54 Terabytes / Experiment / Station

This is a reasonable number for long term storage given that a archival file system is used.

You would need ~ 500 Terabytes of storage per station to run the experiments given you can do phase 1 and phase 2 correlation of the RAW data within 7 hours of it “hitting the sensor”.

What type of Network do you need ?

This is **not just a data network problem
it is a **Computational GRID & Storage Problem as well****

A correlator needs to look at the same “wave” as it traverses through the array. It needs to look at the same time stamped data.

Leave the data at the station and have a network SAN Environment built so that a CPU can read the sample files across all 130 stations to carry out the 3rd level of correlation.

This is best done with a software correlator running as a GRID application as the nature is a very parallel process.

You DON'T need to copy the 500 Terabyte file around, you just need to have each CPU doing the correlation for time stamp “X” read its little bit on all 130 sites and then grind away at the Fourier transform.

What type of Protocols do you need ?

We need to define a new series of protocols that handle parallel high performance streams of secure block data.

The protocol is needed to search and identify datasets in a block mode in a distributed, ordered, secure manner where computing resources can work remotely on the data **without moving** the whole dataset.

Datasets need to be index addressable across a network.

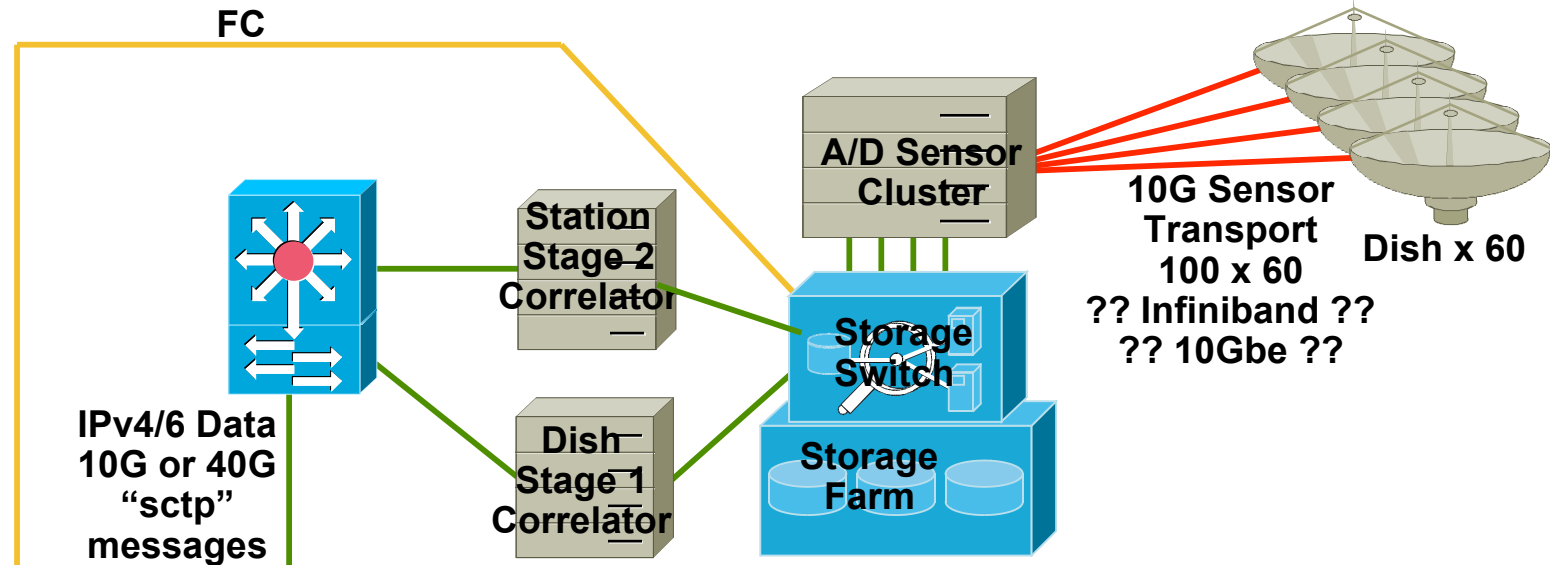
Looking at this in a L4 paradigm we need multiple streams, security, defined messaging & performance.

sctp with a better congestion control algorithm looks kind of right as a starting point...


An SKA Data Architecture

- ??? ???
- 10 or 40 Gigabit Ethernet
- 10 or 40 Gigabit FC

SKA Station Site




SKA Stage 3 Correlator Site



Radio Astronomy Scientific Research is at a Tipping Point

Inspirational leaps forward in Science are now achievable with integrated network enabled technologies

- Within the last 3 years advances in general purpose computing calculation & storage capabilities coupled with high performance, low cost 10GbE networking and high definition sensor resolution have converged to create a tipping point in Radio Astronomy.
- The discipline is moving away from custom-built correlator and computation engines to flexible multi-purpose software systems integrated with networking to bind instruments directly to distributed computation grids and P2P environments.
- The outcomes from this fundamental will undoubtedly result in advances similar in size and scope that the invention of http had in CERN in the 1990's.



Reference Sites

SKA

www.skatelescope.org

CSIRO Australia Telescope National Facility

www.atnf.csiro.au

Swinburne Centre for Astrophysics & Supercomputing

astronomy.swinburne.edu.au

With thanks to Prof. Matt Bailes for content & vision.

Q and A



