Developments in Heterogeneous Computing at Green Bank

John Ford

GUPPI Team: Patrick Brandt, Ron Duplain, Paul Demorest, Randy McCullough, Scott Ransom, Jason Ray

Heterogeneous Computing Systems

A system made up of architecturally diverse user-programmable computing units.

Typically a mixture of two or more of the following:

- Traditional RISC/CISC CPUs
- Graphical Processing Units (GPU)
- Reconfigurable Computing Elements (RCE)

Real-time Computing

A real-time computing system must operate while meeting fixed deadlines for completion of scheduled computations

Computations finished correctly but late are not useful

Real-time does not mean "really fast"!

But it is helpful!

Heterogeneous Real-time Computing

- Combines both ideas into one system
  - Note that many heterogeneous systems are not designed to be real-time systems:
    - SRC-{6,7}, Convey, CPU/GPU clusters, supercomputers
  - Building a heterogeneous system is more difficult than building a homogenous system due to the need to master multiple toolsets, computing models, and hardware characteristics.
  - Add in the requirement for real-time response, and it is even more challenging.

Then Why Bother?

- The data rates are too high for a general-purpose machine
  - Use an FPGA preprocessor to operate on the raw data streams for I/O management
  - Use a GPU to offload computations from the CPU for CPU load management
- A general purpose machine would demand too much power
  - FPGA's and GPU's are more efficient at computations that fit their abilities

Computer Systems Architecture

Processor architectures

- CPU
  - X86_64 – Intel and AMD
  - Everything else – IBM Power and Cell
- GPU
  - NVidia
  - AMD/ATI
- Reconfigurable Computing Elements (FPGA-based processors)

Reconfigurable Computing Elements

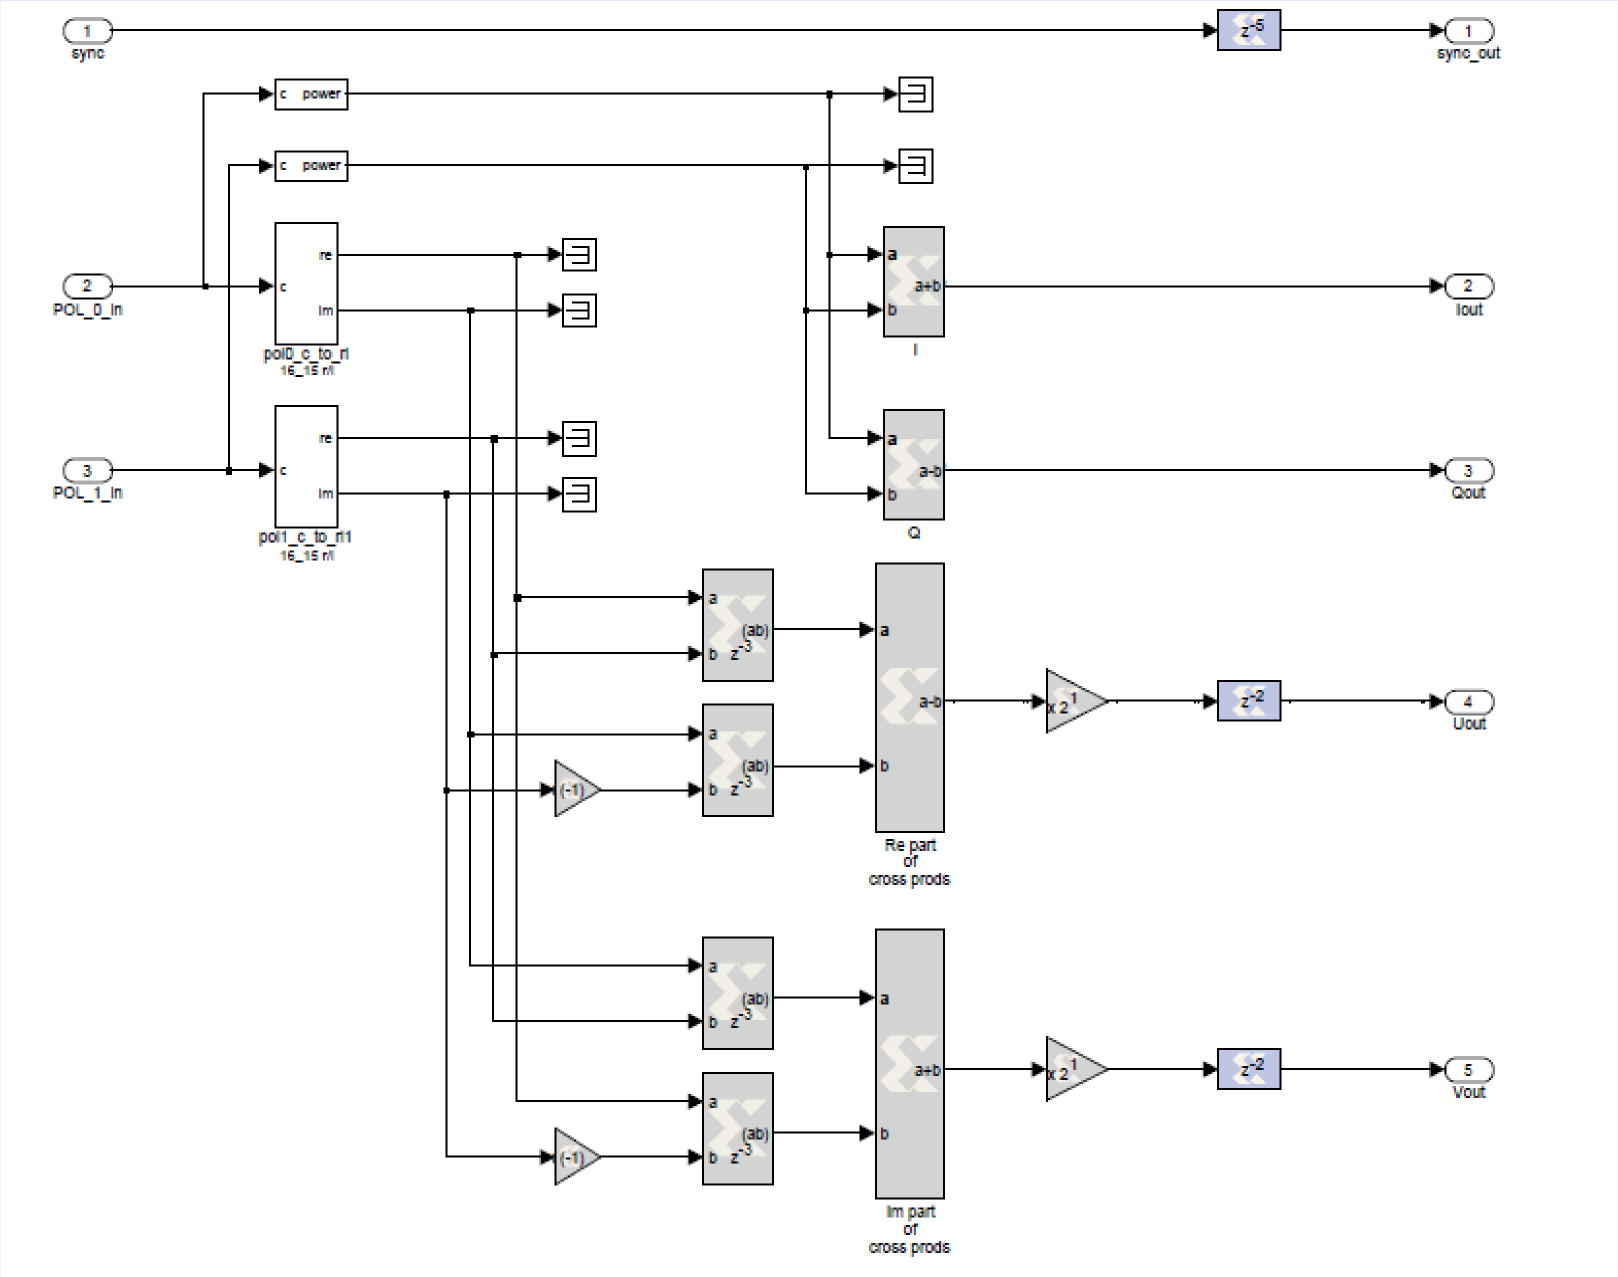

Pulsars

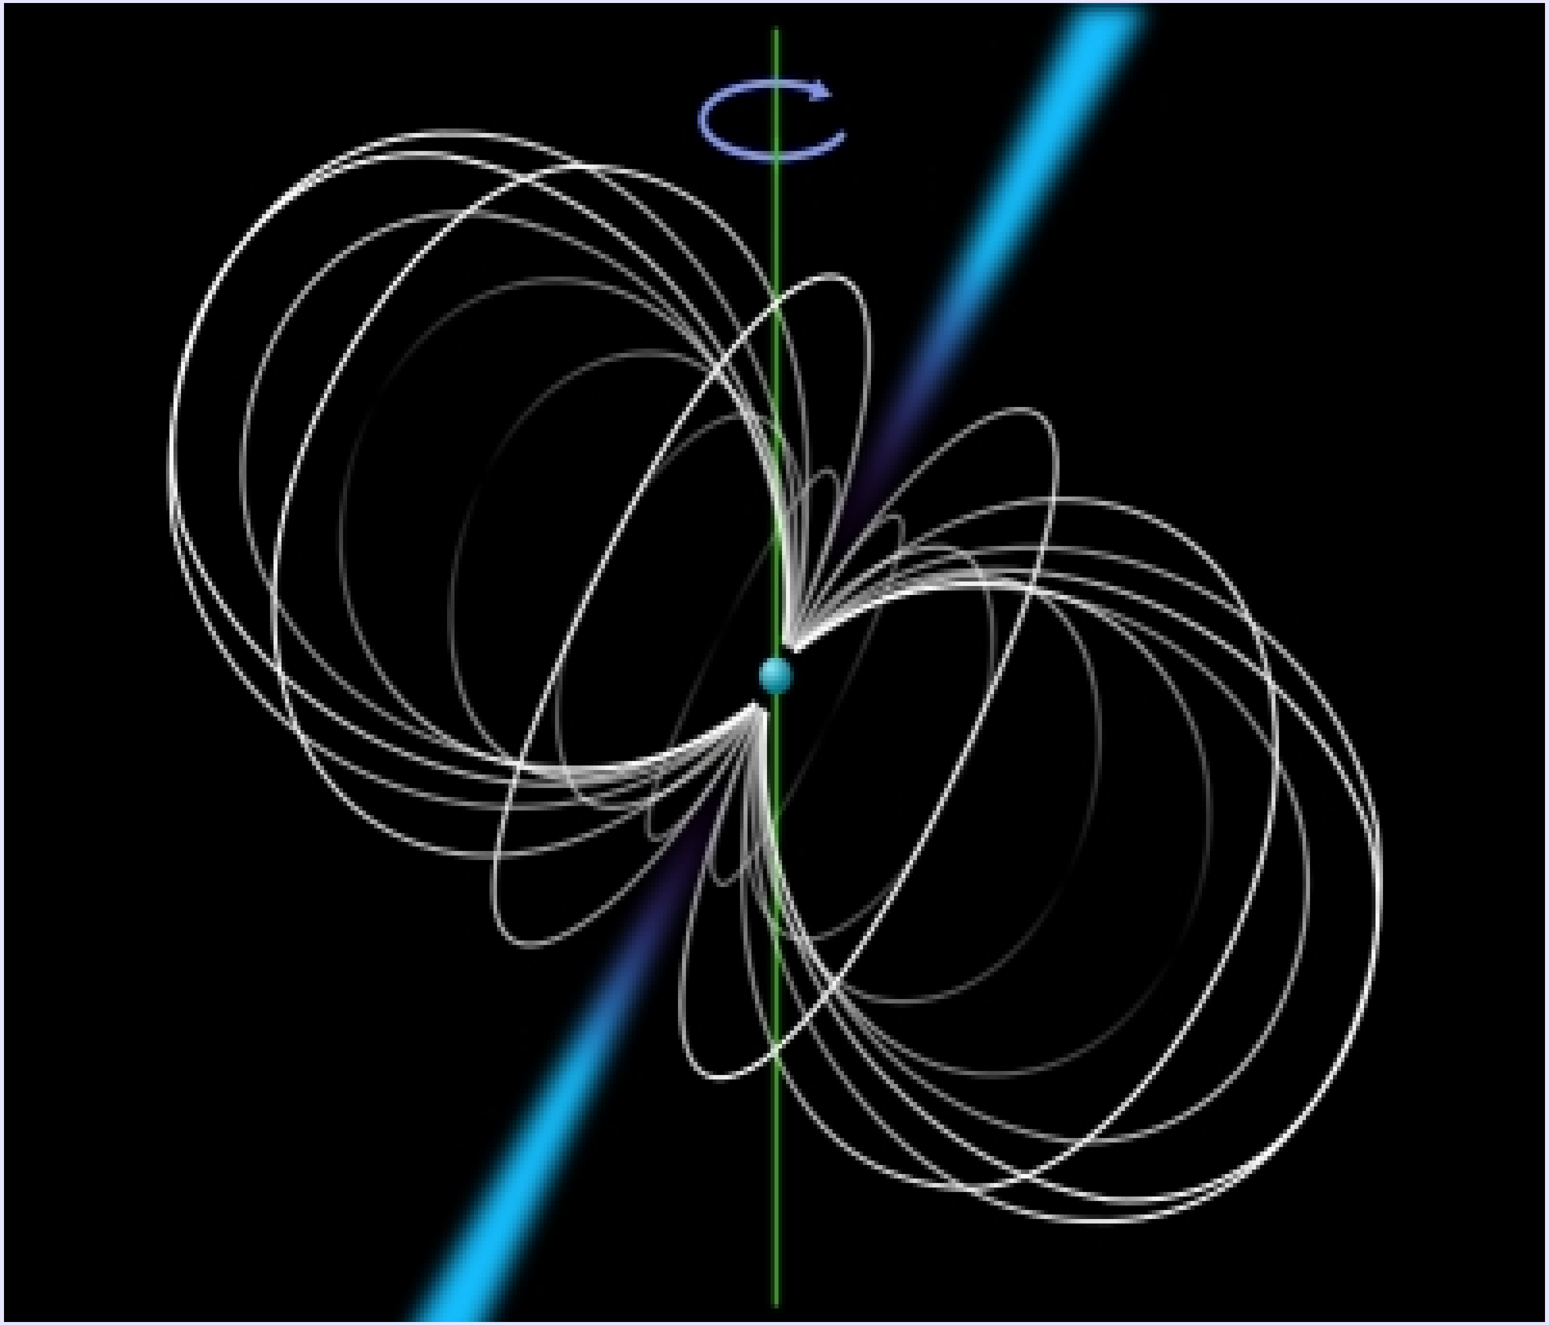

An Example: GUPPI

Pulsar “Search” backends:

- Fast-dump spectrometers
- Time resolution 50 to 100 μs
- Frequency resolution 25 kHz to 1 MHz
- Often trade data quality for more BW

Pulsar “Timing” backends:

- Pulse period folding
- Coherent dedispersion
- High time resolution, ~1 μs
- Data quality (# bits, polns, etc) more important

GUPPI Block Diagram

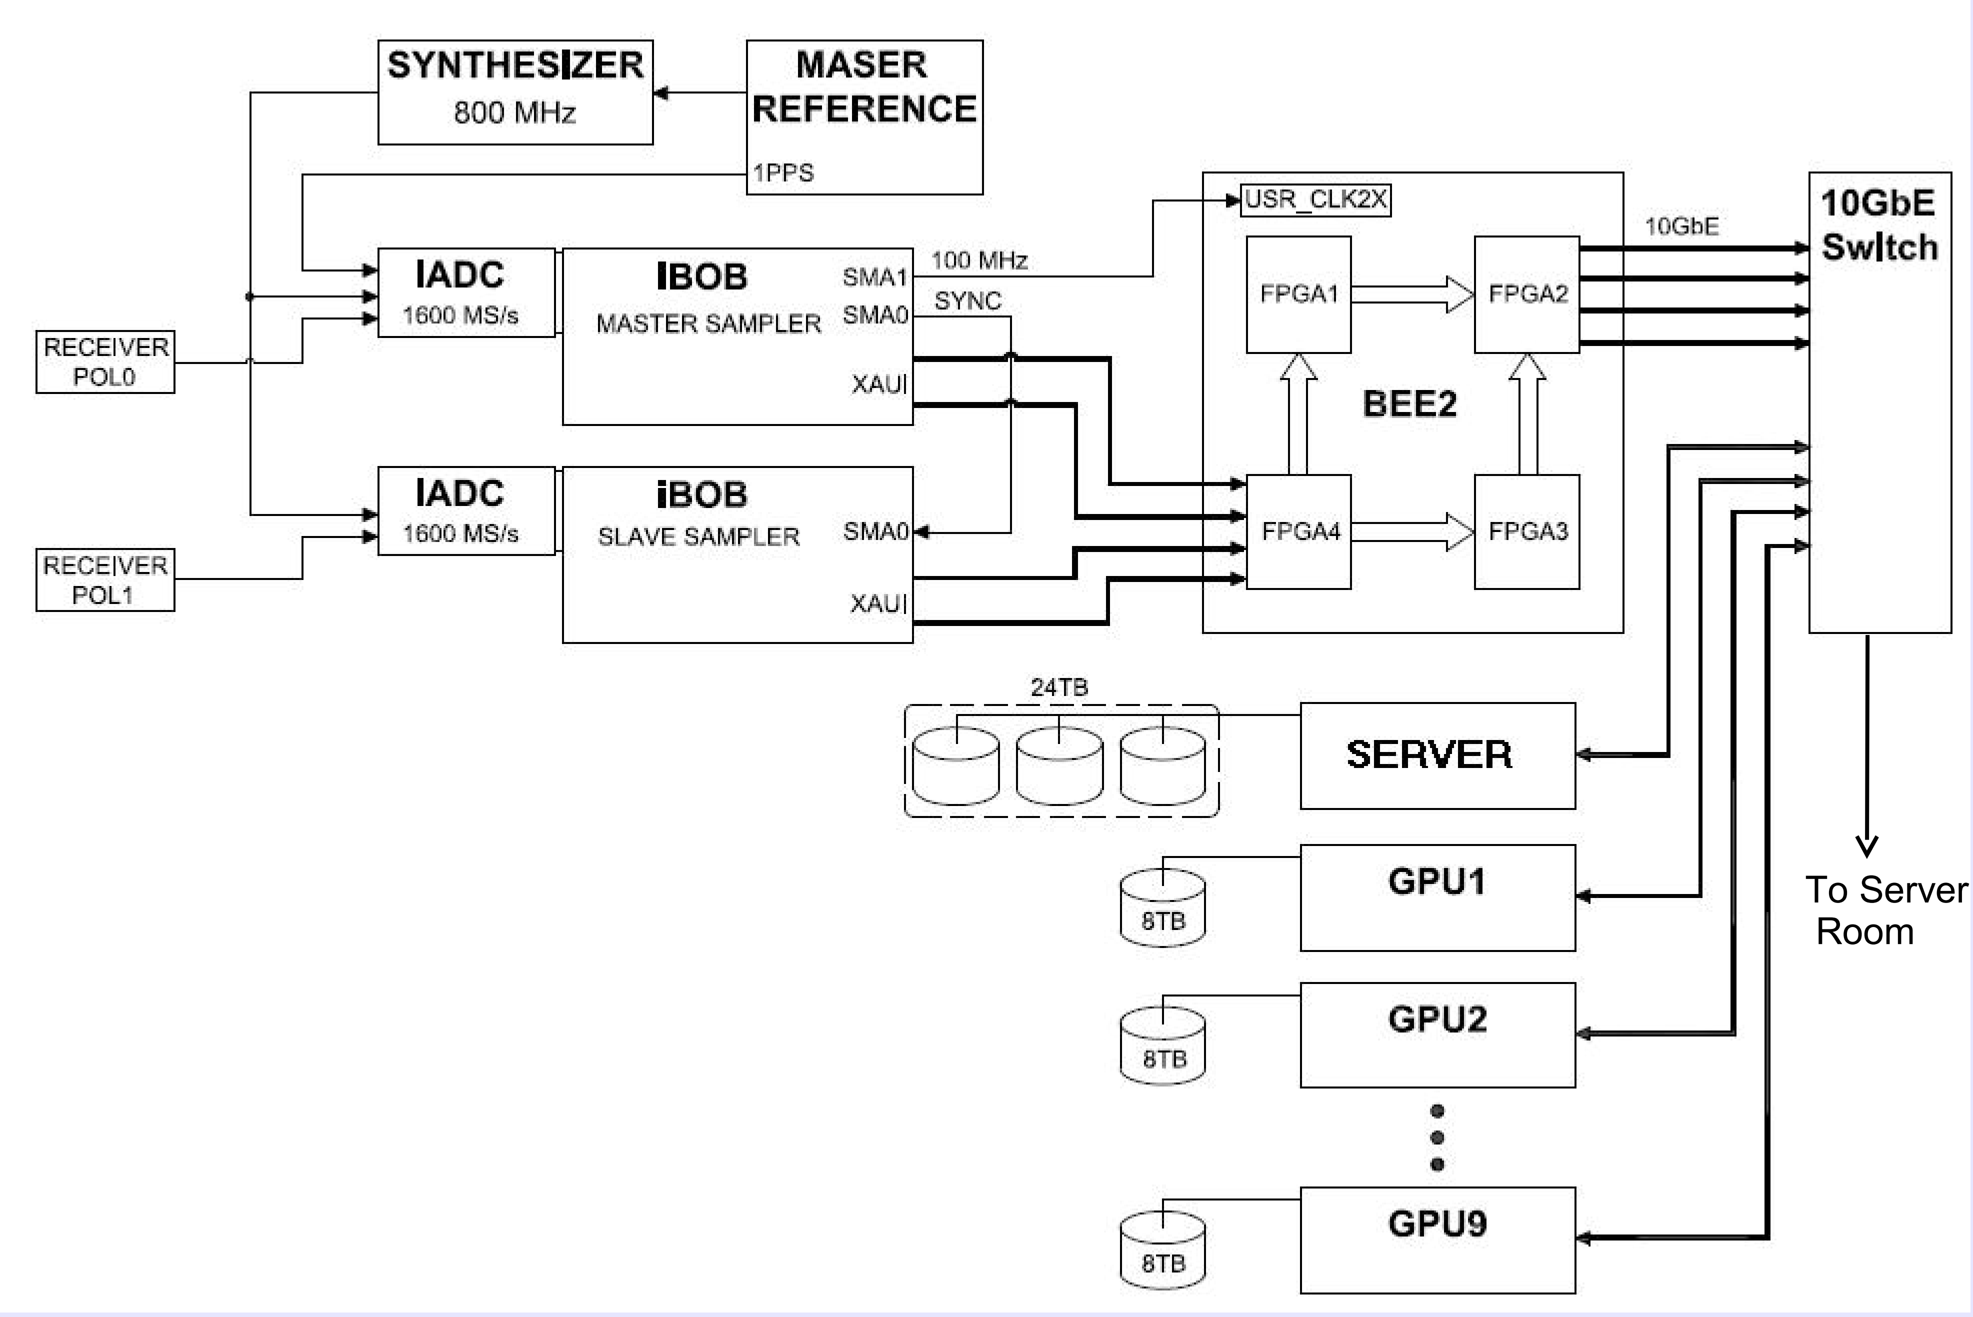

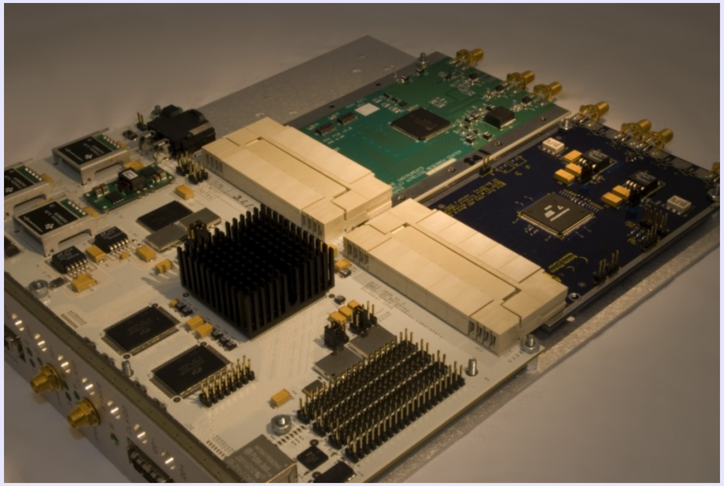

GUPPI architecture:
~1 MHz PFB in FPGAs
Coherent dedisp in GPUs

XAUI

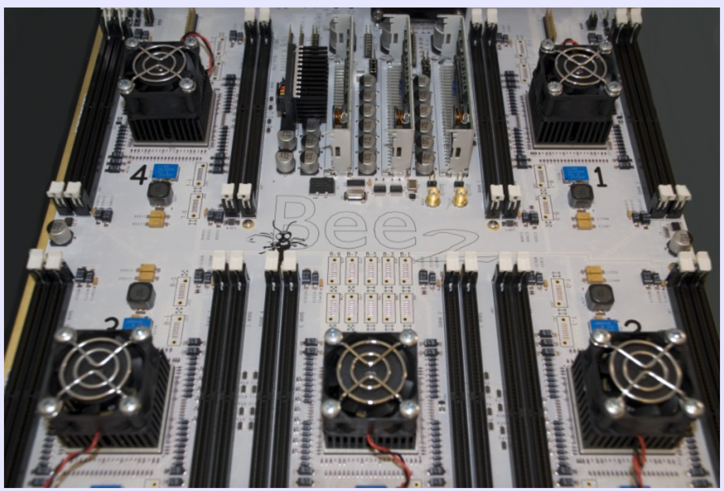

10 Ge
24 Gb/s

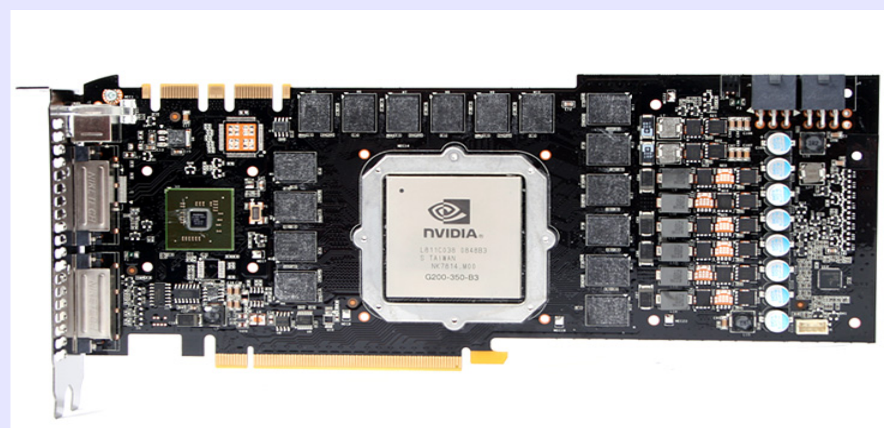

P. Demorest

>10x improvement in BW!

Fully utilizes all GBT low-freq receivers.

Improved S/N ratio, also reduces scintillation variability.

~ 350 MB/sec sustained data rate (~10 TB/8 hours)

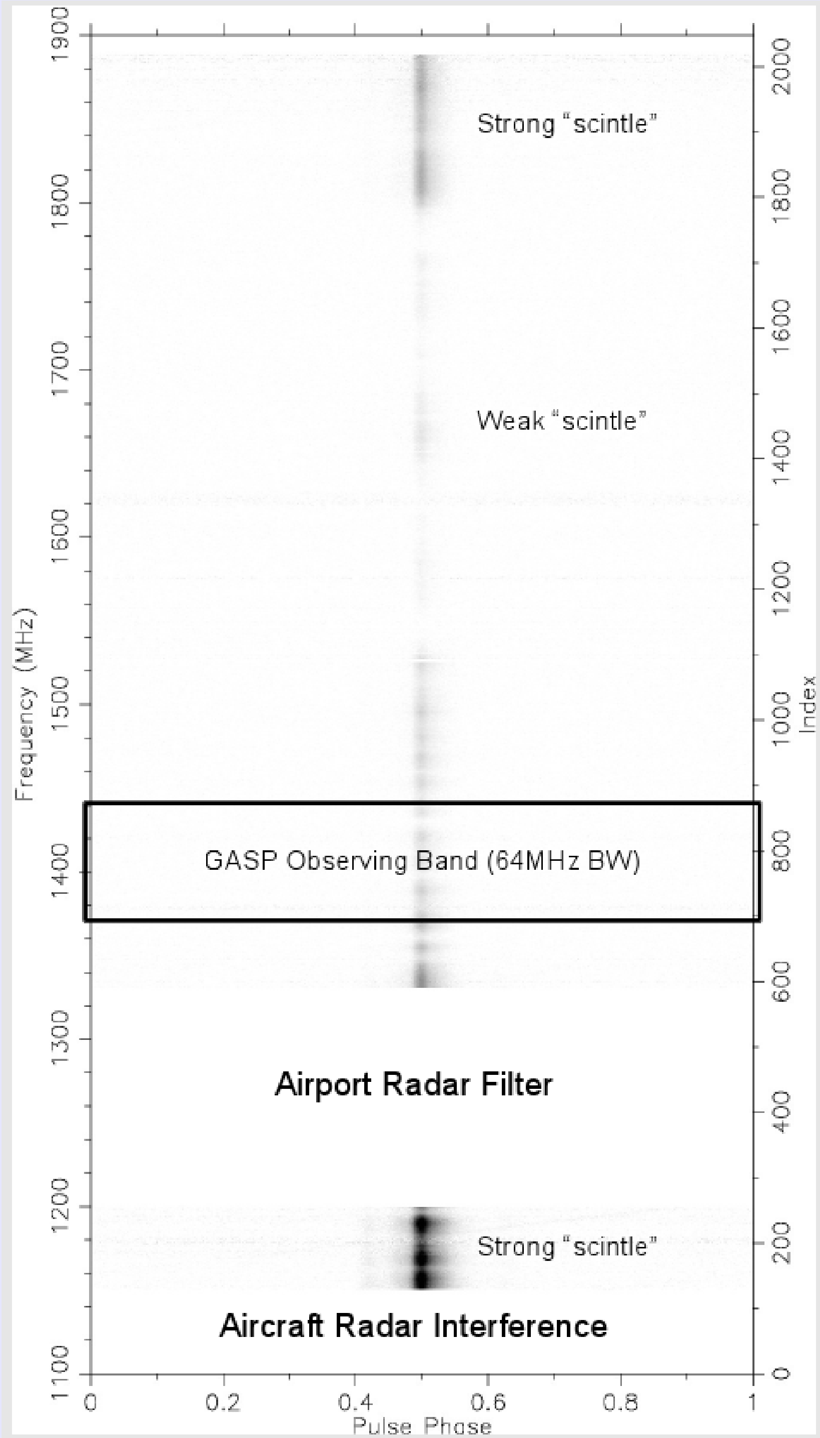

PSR J1713+0747
plot: S. Ransom

Profile comparisons

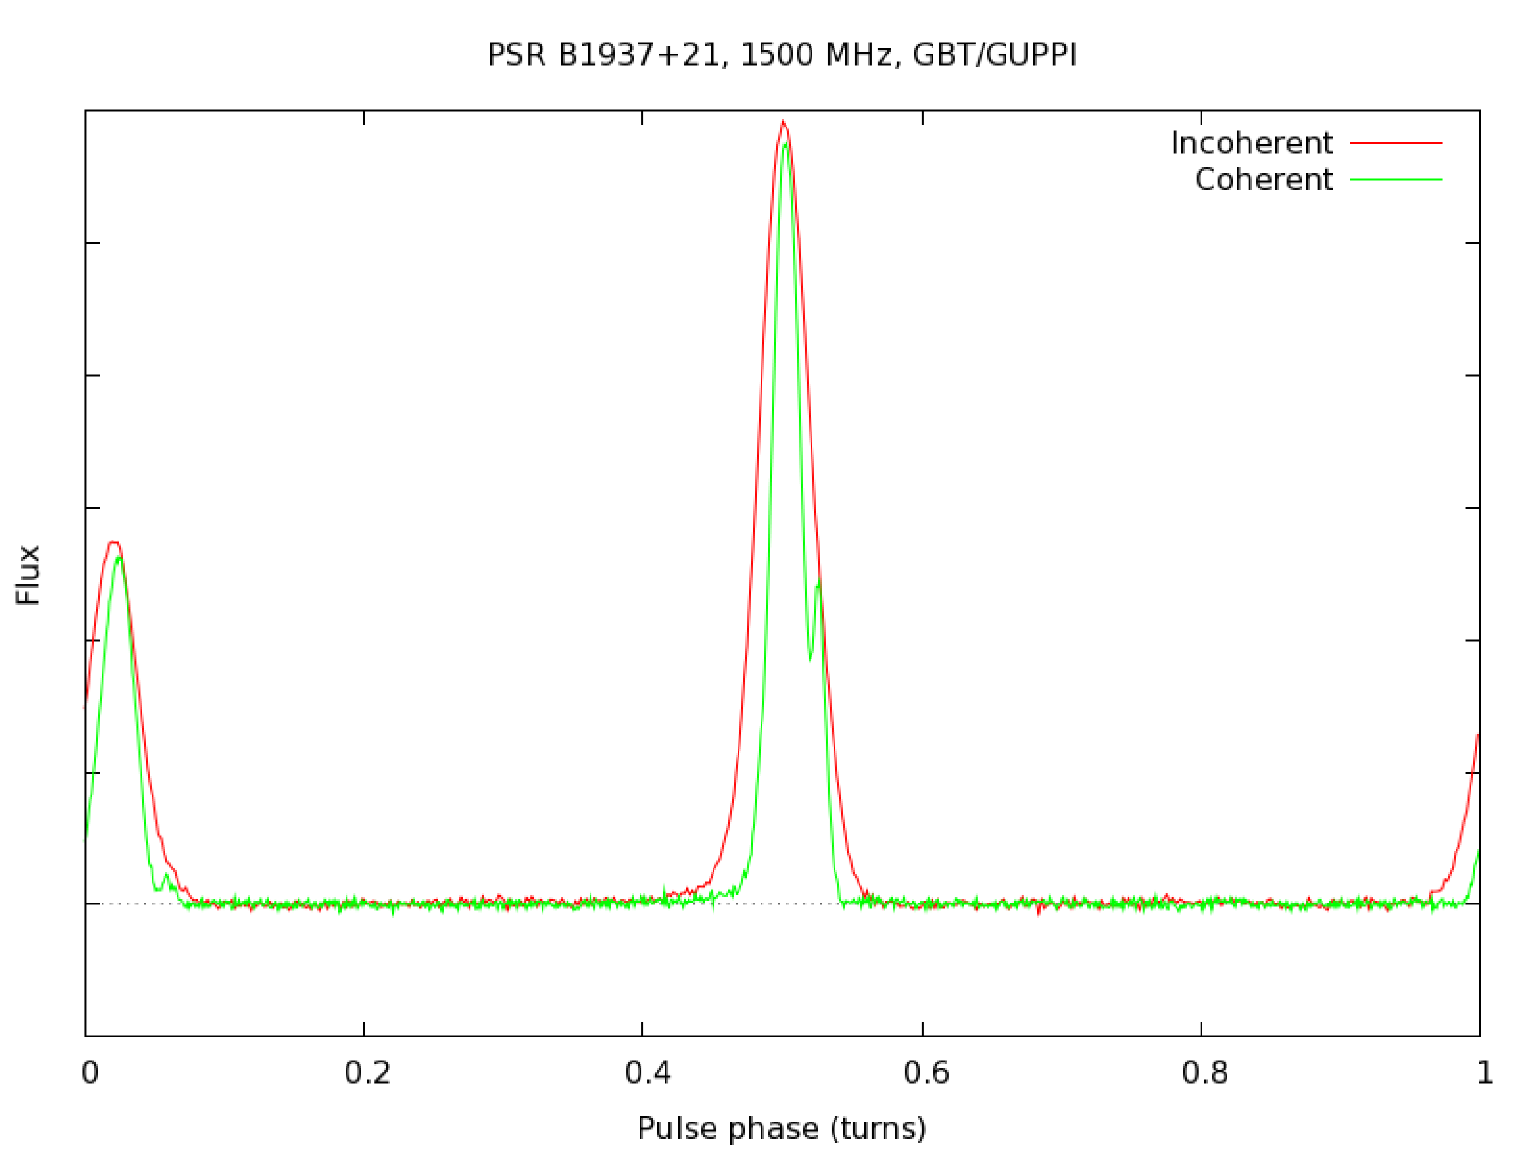

Conclusions

Heterogeneous computing works.

- FPGAs for the fastest streaming integer signal processing
- GPUs for fast floating point DSP calculations
- General-purpose CPUs for managing the system, assembling and transmitting data, and for parts of the computations that are not amenable to the GPU

New GBT Spectrometer

Joint effort with UC Berkeley – NSF ATI project

16 IF inputs

- 8 dual-pol beams, or 16 single-pol beams

3 GS/s sampling rate @ 8 bits/sample

Can be ganged to achieve ~10 GHz instantaneous bandwidth on 2 polarizations

GBT Spectrometer

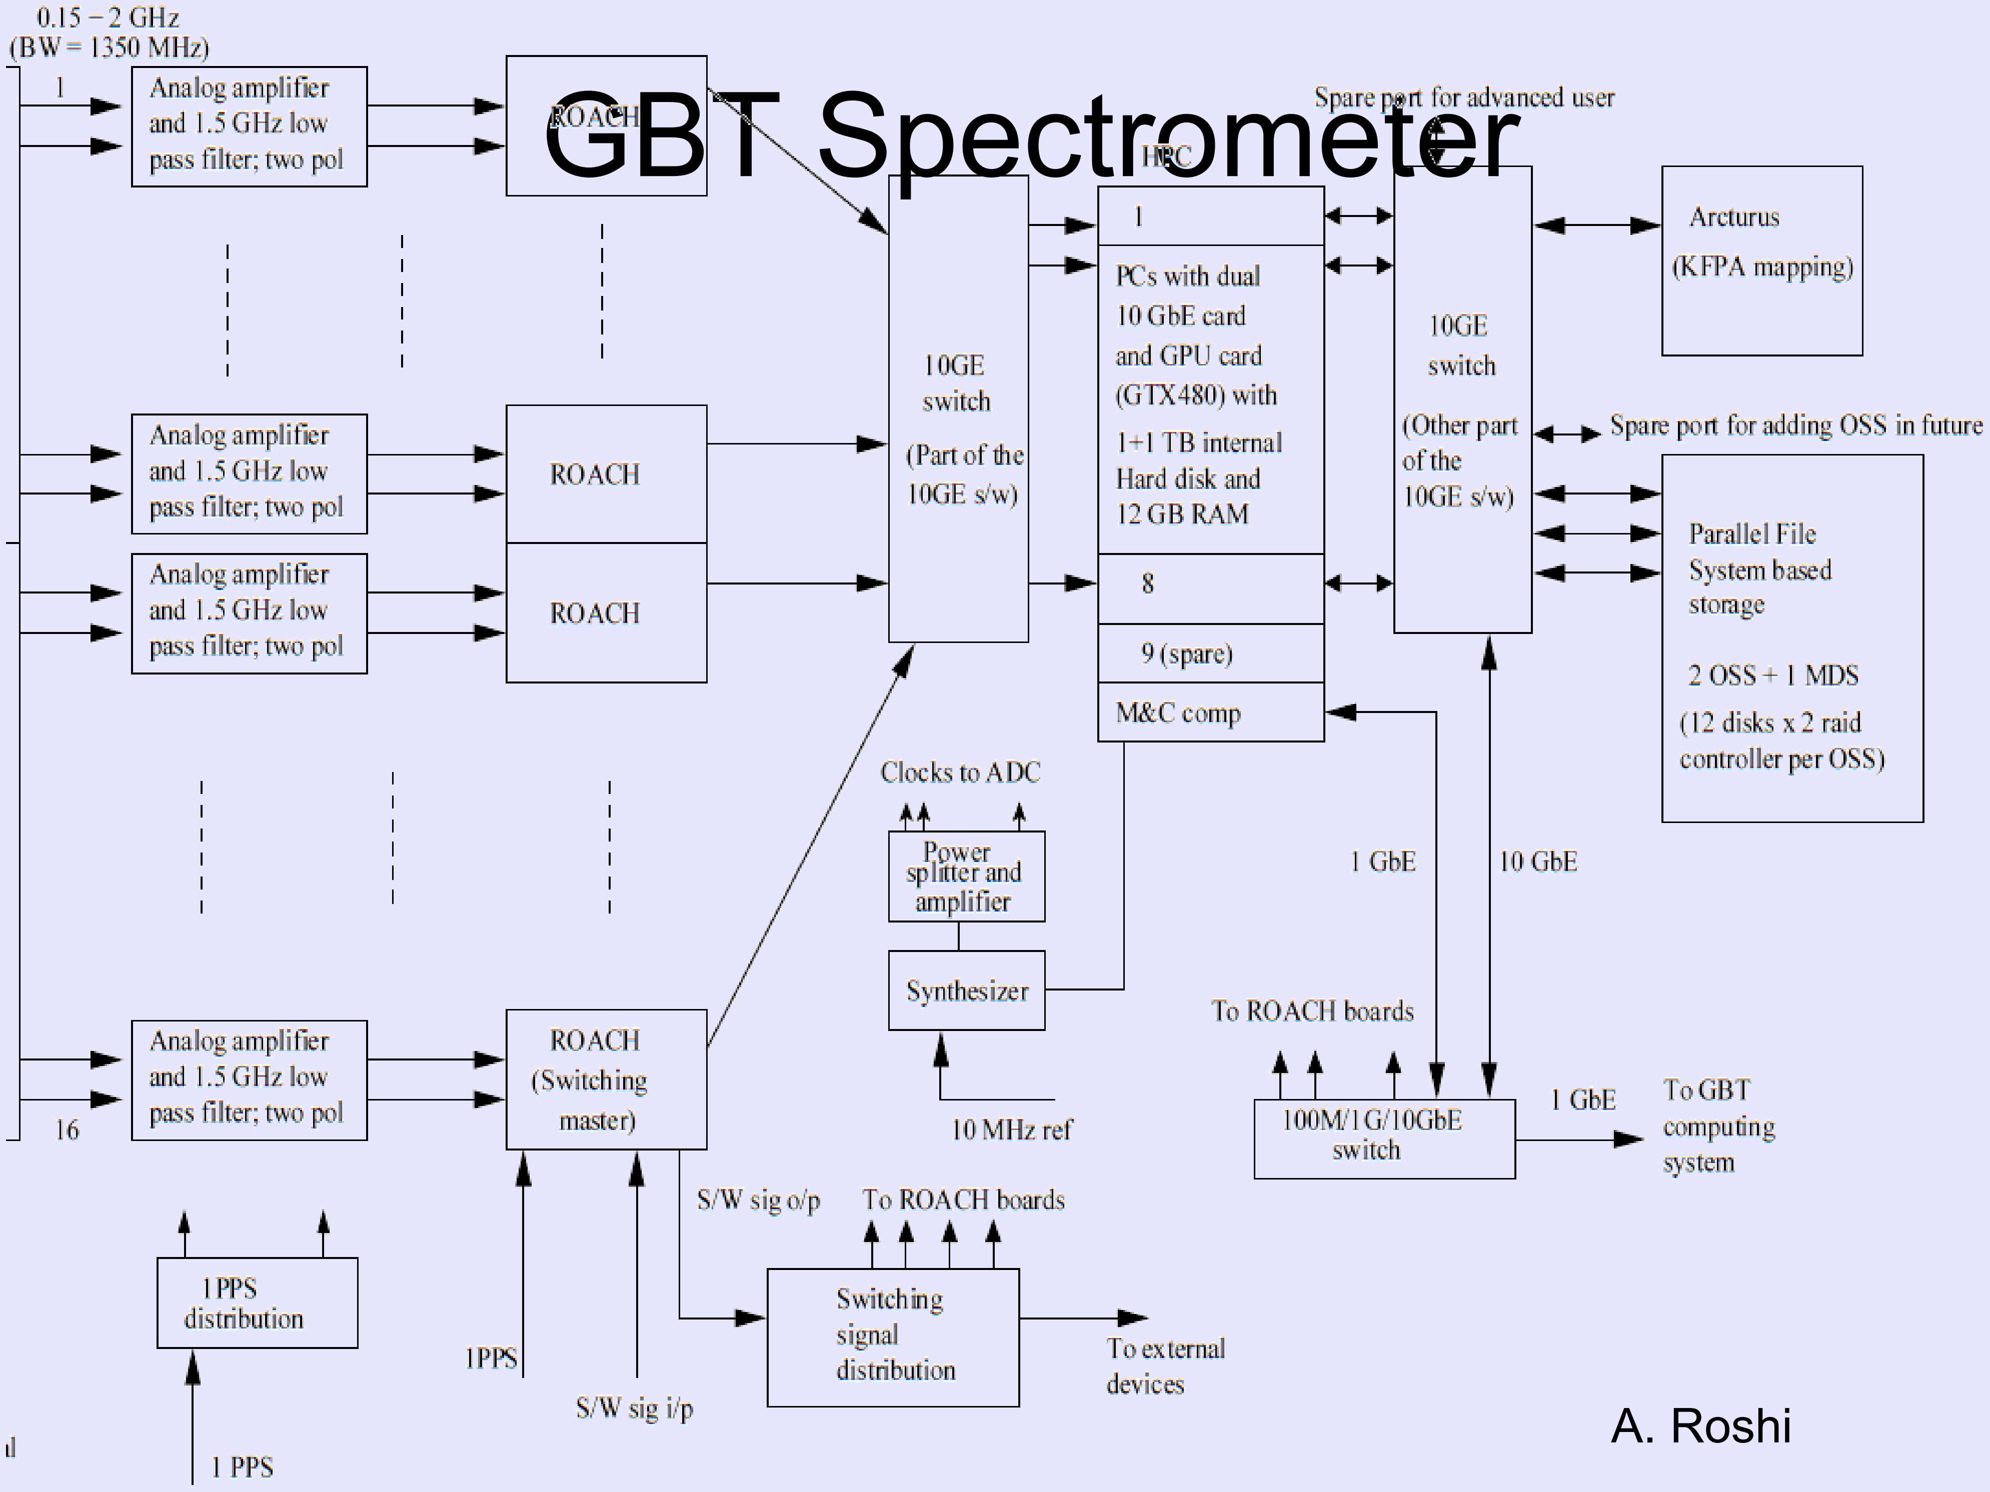

A. Roshi

Spectrometer Data Rates

ADC to FPGA: 3GS/s * 16 Polns * 8 bits = 384 Gb/sec

FPGA to GPU: ~ 9 Gb/sec * 8 GPUs = 72 Gb/s

Raw (mostly) Data Output Rate:

Maximum specified at 33 MB/s * 16 polns = 525 MB/s, or 15 TB/8 hours, or 1 PB every 66 hours.

Front-end (FPGA-GPU) hardware is capable of ~10x this rate…

Conclusions

Even today:

We can collect data faster than we can write it to disk

We can collect data faster than humans can monitor it for quality

We can’t afford to store raw data even @ the raw disk drive price of $90K/PB.

Our RF bandwidth is >4x our digitized bandwidth

Sign in to the AWS Management Console | Create an AWS Account | English

Search: AWS Product Information

AWS | Products | Developers | Community | Support | Account

Build a Cluster Computing Environment in under 10 minutes.

Try free* with a $20 credit

High Performance Computing | Why AWS

High Performance Computing (HPC)

Researchers and businesses alike have complex computational workloads such as tightly coupled parallel processes or demanding network-bound applications, from genome sequencing to financial modeling. Regardless of the application, one major issue affects them both: procuring and provisioning machines. In typical cluster environments, there is a long queue to access machines, and purchasing dedicated, purpose-built hardware takes time and considerable upfront investment. With Amazon Web Services businesses and researchers can easily fulfill their high performance computational requirements with the added benefit of ad-hoc provisioning and pay-as-you-go pricing.

Follow along with this video tutorial to set up your HPC environment in less than 10 minutes. Please ensure that you terminate your instances after completing the tutorial.

Get Free Credit

How to Get Started

1. **Register for free credit**
2. **Sign up for Amazon EC2** (Amazon Elastic Compute Cloud)
3. **Follow video tutorial**
4. **Terminate your instances** after completing the tutorial

Summary of Required Observing

Table 2: GBT spectrometer modes specified per beam (2 IFs)[a]

Number of sub-bands per IF	Sub-band Bandwidth[b] (MHz)	Number of channels per sub-band per IF	Spectral resolution (KHz)	Velocity range at 90 GHz (km s^{-1})	Velocity resolution at 90 GHz (km s^{-1})	Integration time minimum (msec)	Integration time maximum (sec)
Observing Mode 1							
1	1500[c]	1024	1465	5000[b]	4.9	0.5	60
1	1000	2048	488	3333	1.6	0.7	60
1	800	4096	195	2667	0.7	1.3	60
1	500	8192	61	1668	0.2	2.5	60
1	400	16384	24	1333	0.08	5	60
1	250	32768	7.6	833	0.03	10	60
1	100	32768	3.1	333	0.01	10	60
1	50	32768	1.5	166	0.005	10	60
1	25	32768	0.8	83	0.003	10	60
1	10	32768	0.3	33	0.001	10	60
1	5	32768	0.15	17	0.0005	10	60
1	1	32768	0.03	3	0.0001	10	60
Observing Mode 2							
8	30	4096	7.3	100	0.02	10	60
8	15	4096	3.7	50	0.01	10	60
8	10	4096	2.4	33	0.008	10	60
8	5	4096	1.2	17	0.004	10	60
8	1	4096	0.2	3	0.0008	10	60

[a] These modes are implemented in each spectrometer that processes 2 IFs from a beam.
[b] In Observing Mode 1, bandwidths less than 1500 MHz should be centered between 150 MHz and 1350 MHz.
[c] The usable bandwidth will be 1250 MHz, which corresponds to a velocity range of 4165 km s^{-1} at 90 GHz.

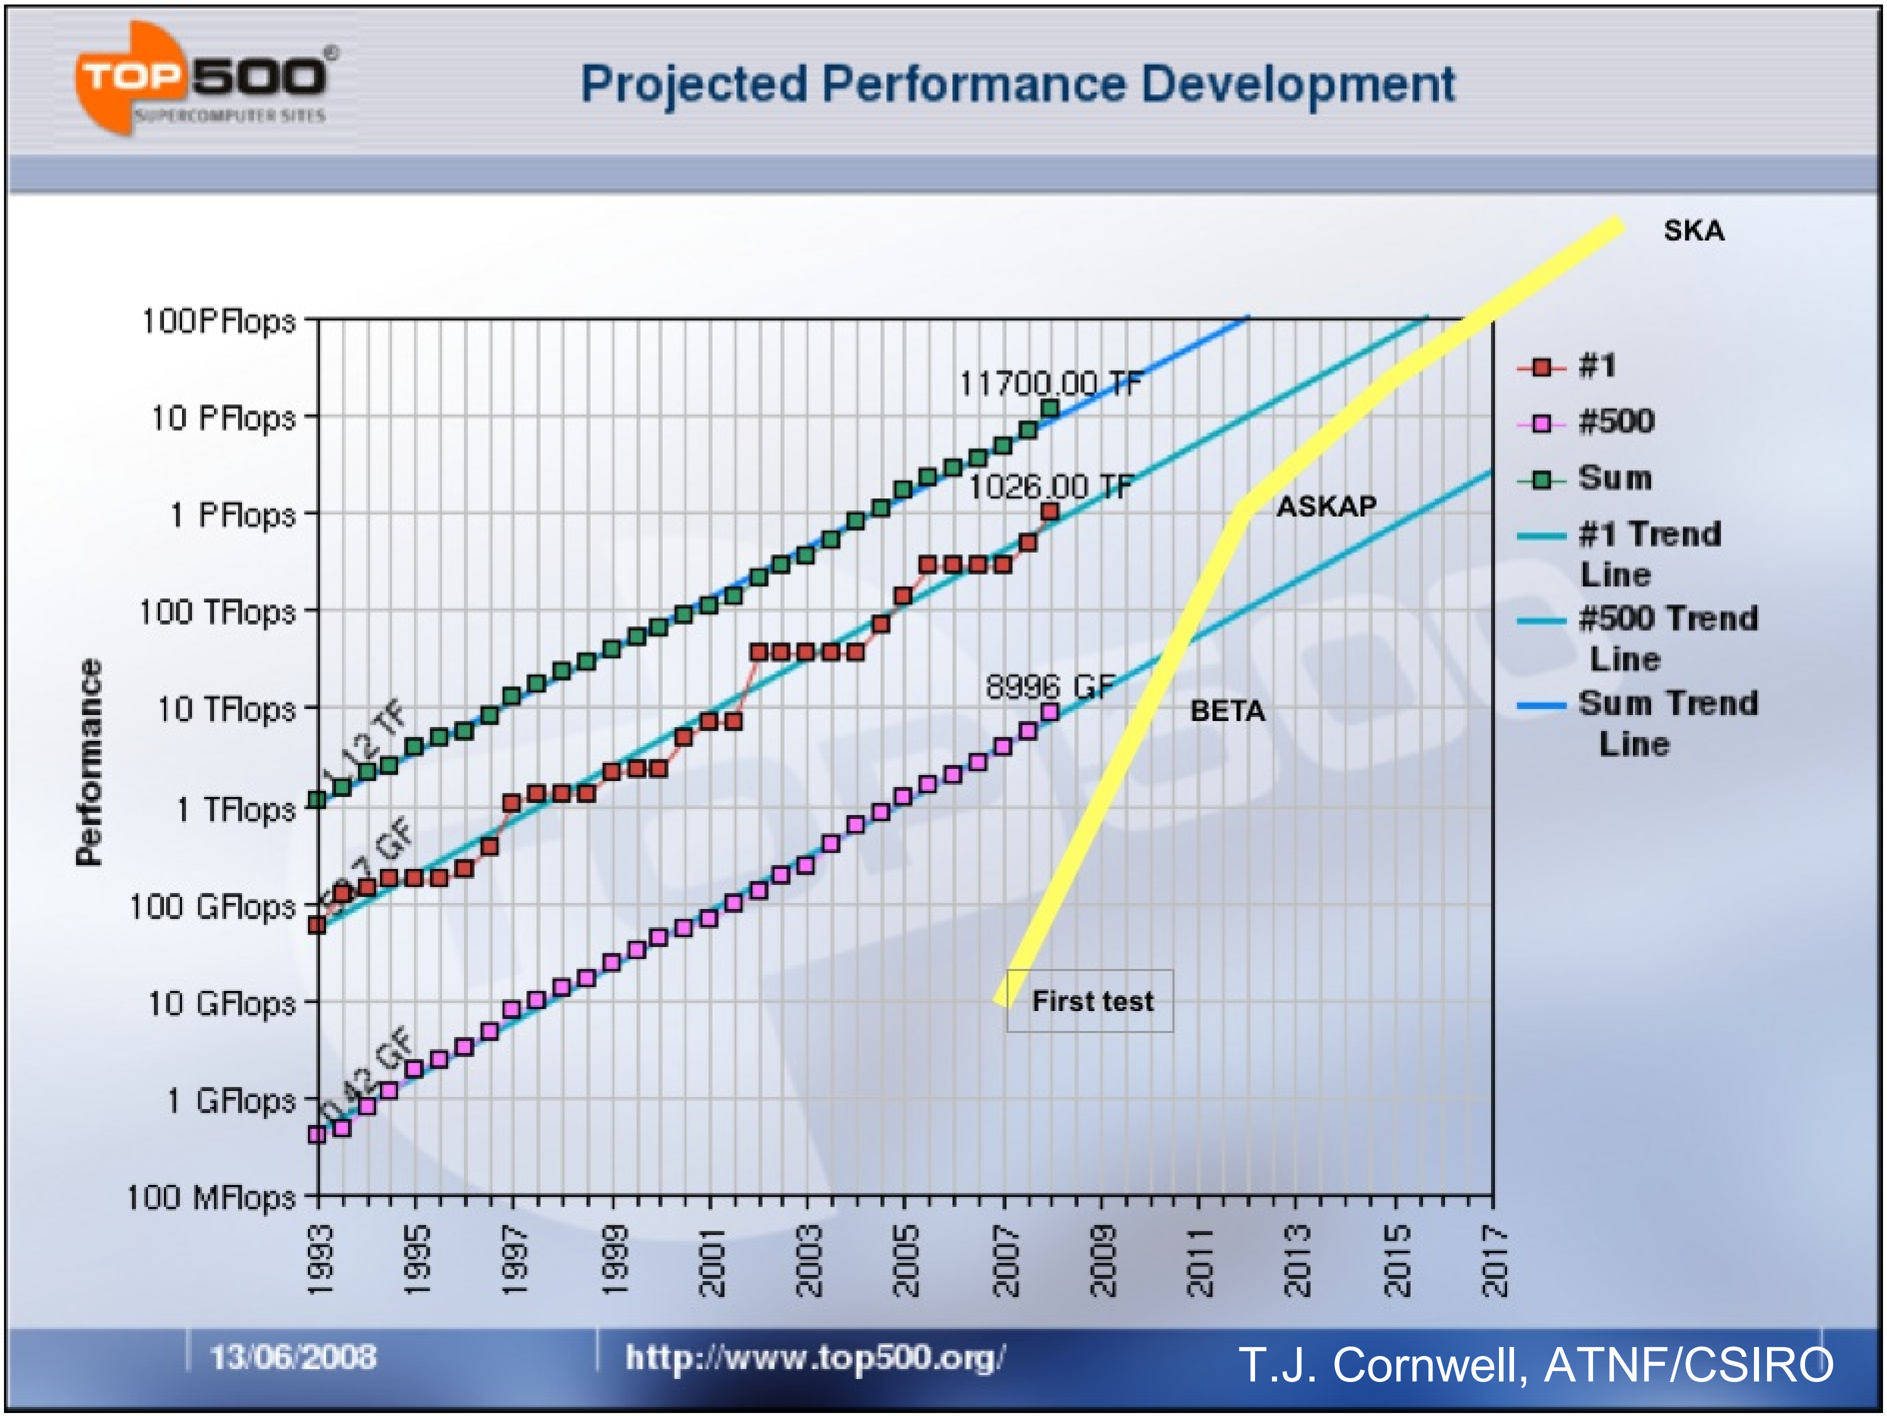
TOP500 SUPERCOMPUTER SITES
Projected Performance Development
SKA
ASKAP
BETA
First test
100PFlops
10 PFlops
1 PFlops
100 TFlops
10 TFlops
1 TFlops
100 GFlops
10 GFlops
1 GFlops
100 MFlops
Performance
11700.00 TF
1026.00 TF
8996 GF
1.12 TF
59.7 GF
0.42 GF
1993
1995
1997
1999
2001
2003
2005
2007
2009
2011
2013
2015
2017
#1
#500
Sum
#1 Trend Line
#500 Trend Line
Sum Trend Line
13/06/2008
http://www.top500.org/
T.J. Cornwell, ATNF/CSIRO

Exa-problems for SKA computing

- Timescale - 2014 +
  - 100% science observations ~ 2020
- Management of complexity
  - Many stake-holders and users
- Networking
  - ~ 1Tb/s per antennas short haul (< 200km)
  - 160Gb/s per station long haul (> 200km)
- Data flow
  - Data rate ~ 1 EB/day
  - Real time processing vital!
- Parallel processing
  - 1 EFlop/s ~ 1 billion cores
  - How do we program these?
- Computing power requirements
  - 1Eflop/s = 350MW at Blue Gene efficiency
  - Need one to two orders of magnitude improvement
  - Plus green power
- Minimize use of radio-astronomy specific technology

Tim Cornwell, ATNF/CSIRO