

e-VLBI with the SFXC software correlator

Mark Kettenis, Aard Keimpema, Des Small,
Mike Sipior, JIVE
Damien Marchal, University of Amsterdam



Netherlands Organisation for Scientific Research



UNIVERSITY OF AMSTERDAM



History

C++ implementation of the SFXC correlator used for the Huygens probe tracking.

Super FX Correlator algorithm developed by Sergei Pogrebenko.

- FX correlator
- Narrow-band, high spectral resolution
- Fractional bit-shift before fringe rotation
- All baselines calculated on a single “node”

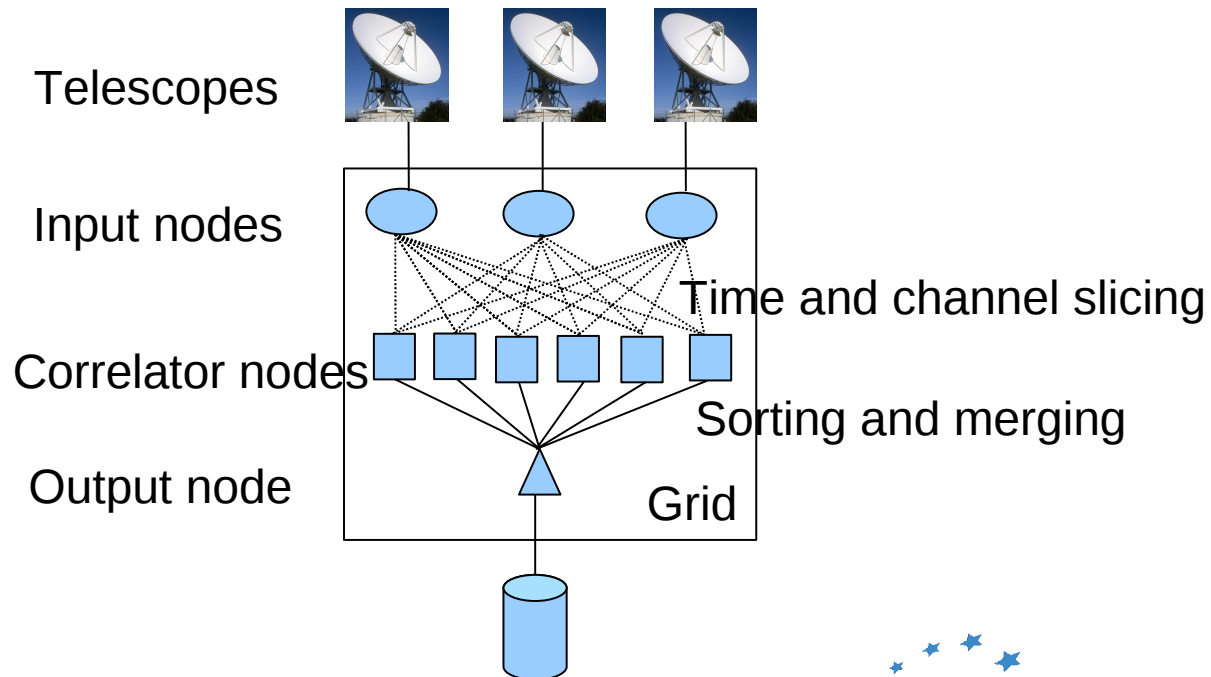
- Largely unoptimized
- Correlator core only



Design Decisions

- Parallelisation in both subbands and time.
- All baselines for a single subband processed on a single cluster node.
- Further course grained parallelisation (in time only) to distribute over multiple clusters.

- Input node (one per station)
- Correlator node (as many as feasible)
- Output node
- Manager node

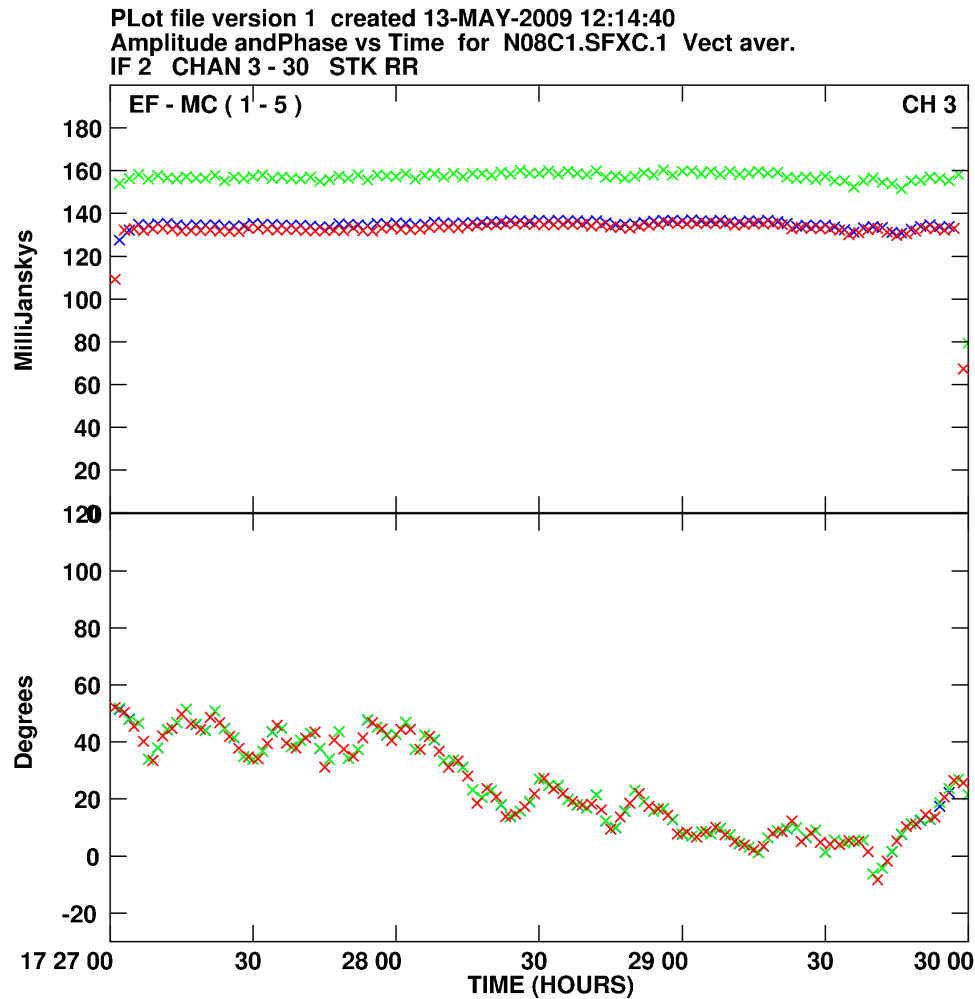


Current State

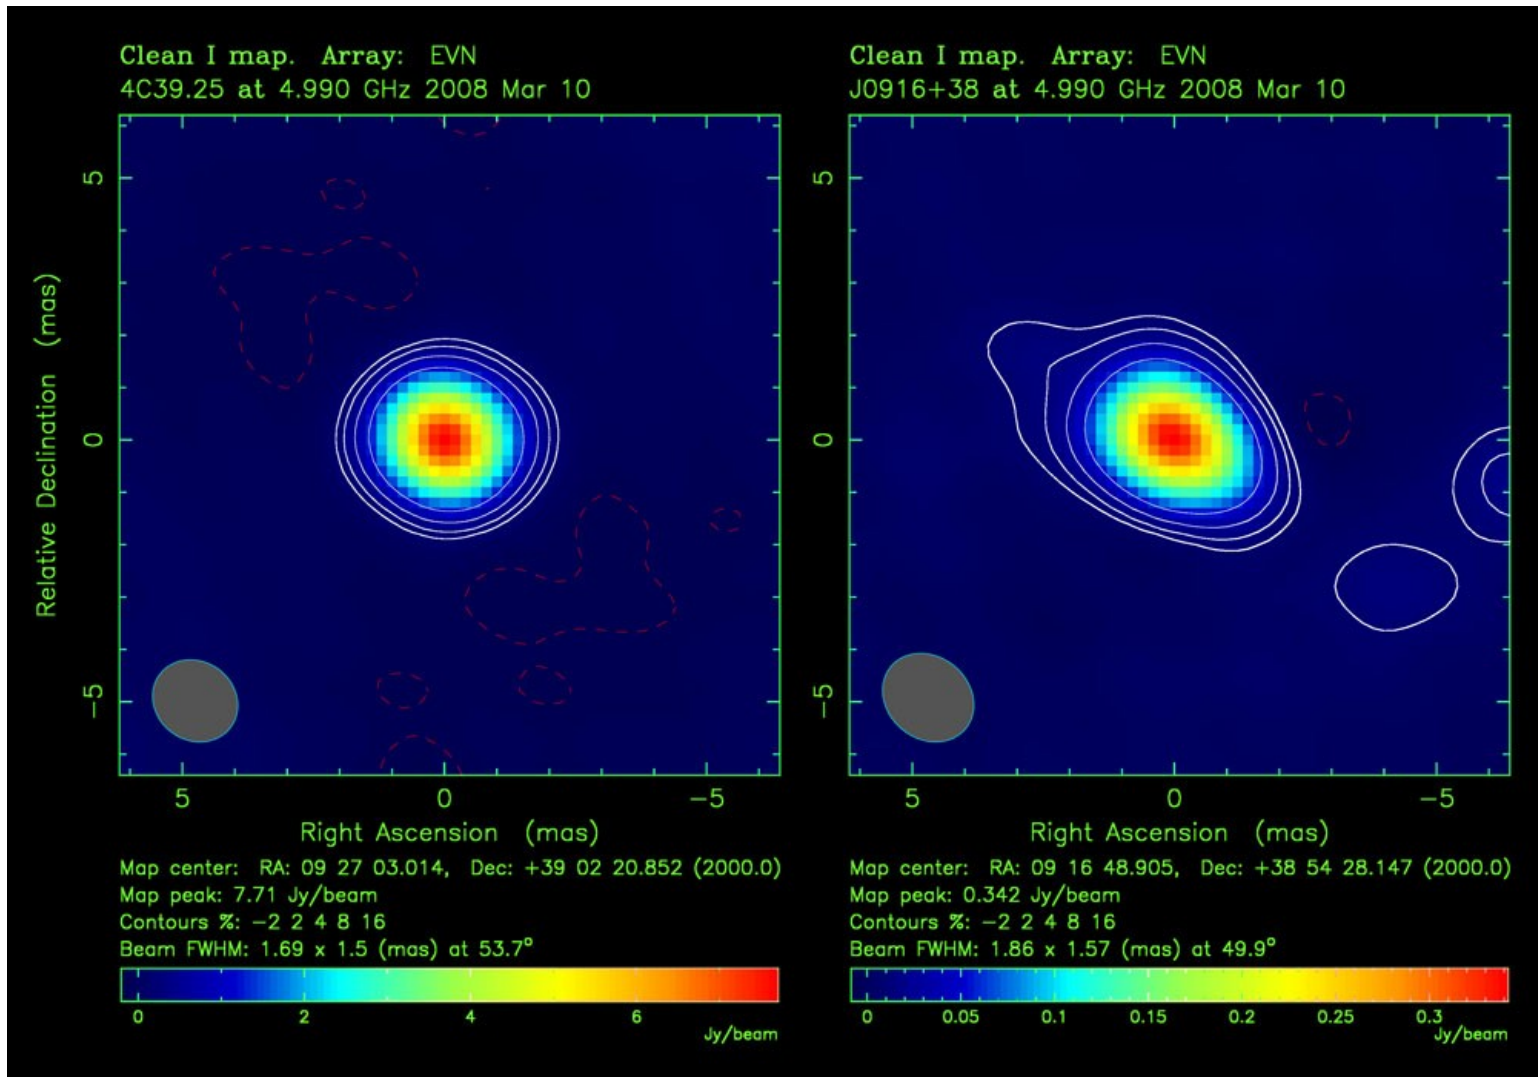
- VEX-driven, with small JSON-based control file specifying correlation parameters.
- Modular: input, correlation and output nodes.
- Parallelized using MPI; distributes subbands and time slices.
- Scales from SMP machines to largish clusters.
- Integrated delay model based on CALC 10.
- All subbands processed, both LSB and USB, cross-correlations.
- Takes Mark4/VLBA (Mark5A) and Mark5B input data.
- Output data can be converted into aips++/CASA MeasurementSet
- Supports both 1-bit and 2-bit samples (and combinations)
- MeasurementSet → FITS IDI using existing tools
- Huge performance increase over the last year!

Validation

- Comparison with the Mk4 Hardware Correlator & DiFX

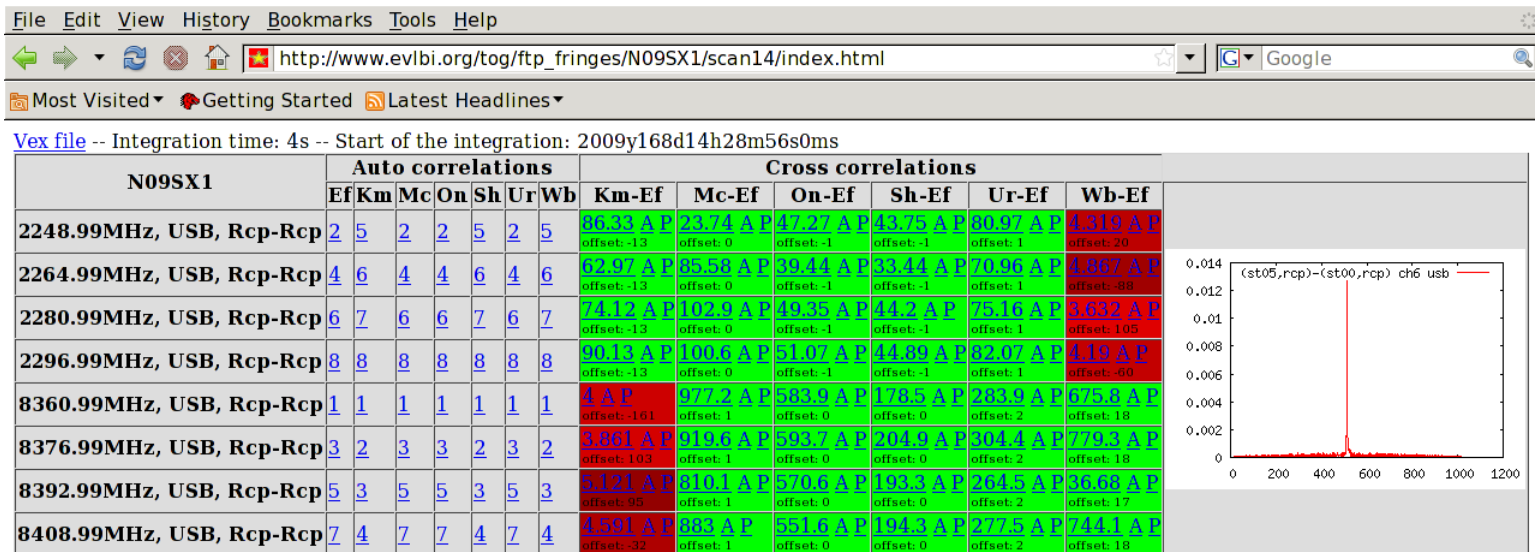


First Image



FTP Fringe Tests

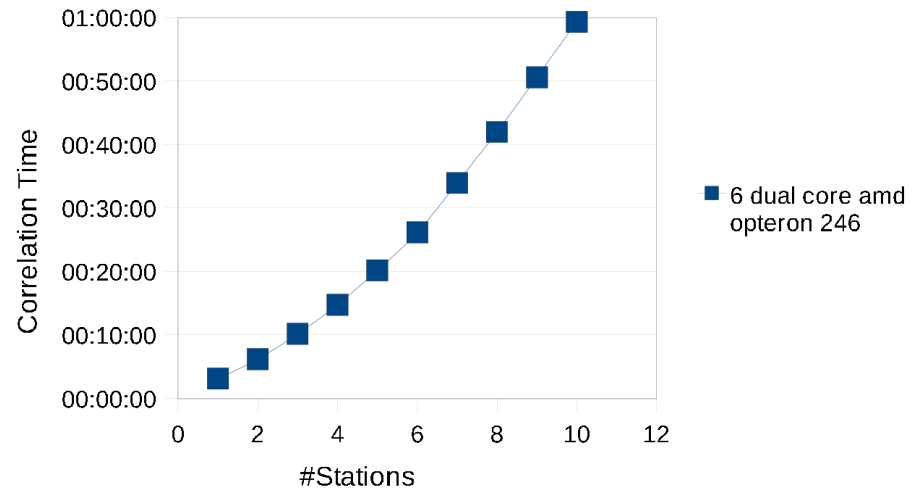
- SFXC used for EVN operation since May 2007
- Interactive web page
- Many changes based on station feedback



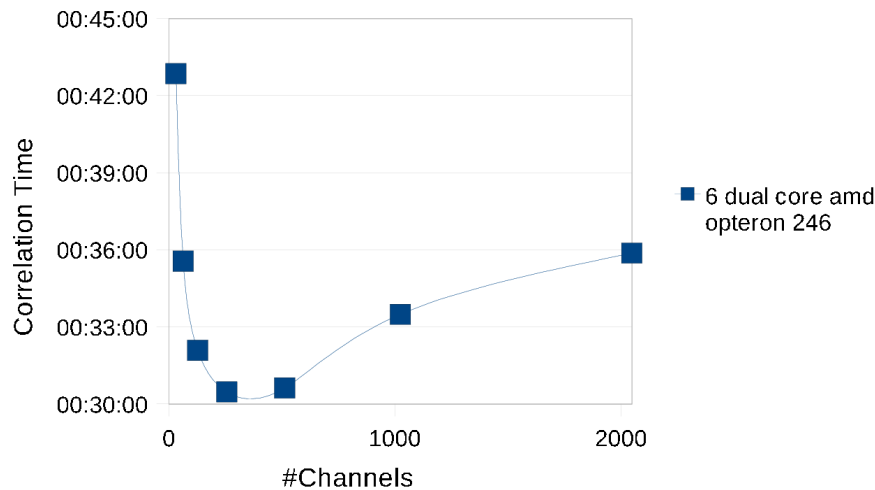
Performance

- Benchmarks done on 6-node cluster
- Dual Opteron 246 machines (2 GHz)
- Performance comparable with DiFX

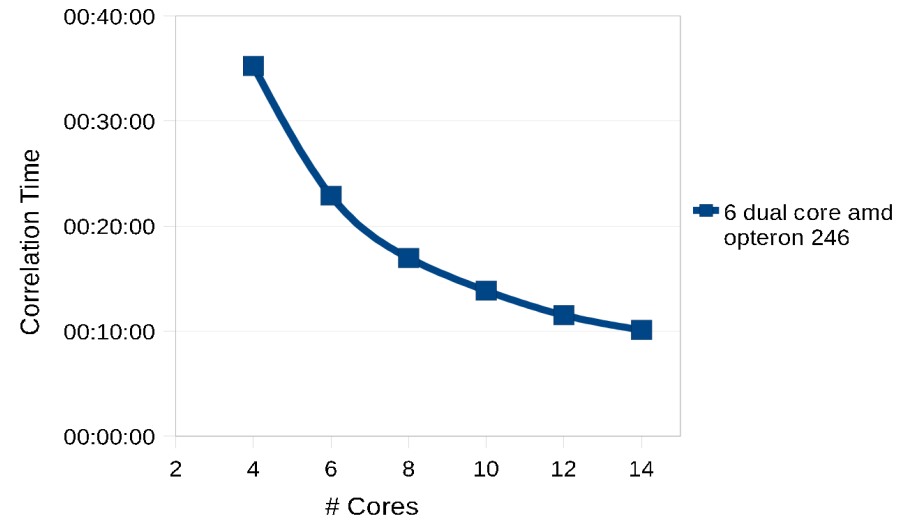
3 minutes data @256 Mb/s



3 minutes data @256 Mb/s (7 stations)



1 Minute data @256Mb/s (7 Stations)



Correlation on GPUs?

- Initial attempt: offload FFTs on GPU:
Improvement less than a factor of 2.

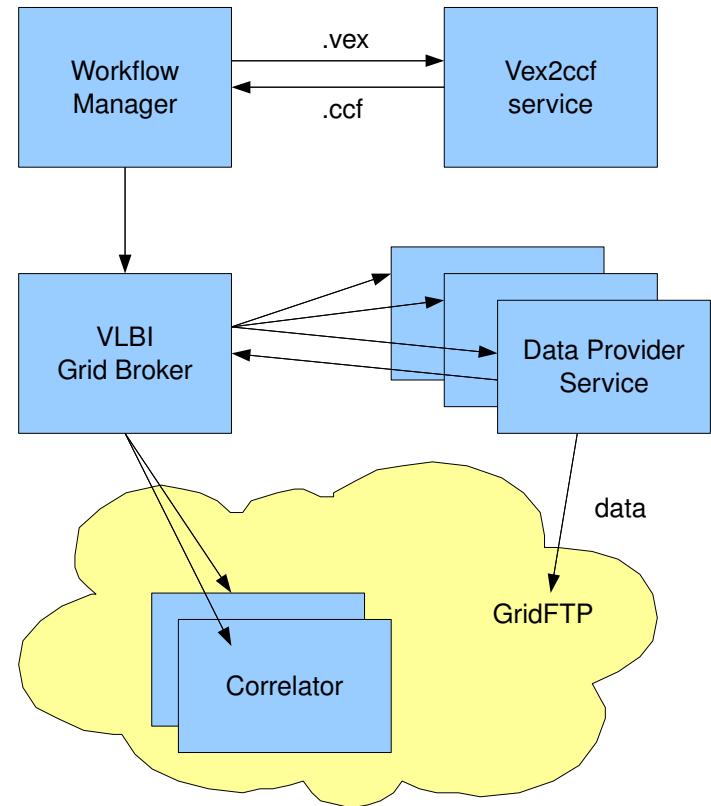
Too much data transfer between main memory and GPU?

- Second attempt: implement full delay compensation and correlation on GPU:
Using CUDA, work in progress.
- NVIDIA Tesla C870 with 1.5GB of memory

Distributed Correlation

Running a single correlation on multiple clusters

- Workflow Manager (PSNC)
- VLBI Grid broker (PSNC)
- Web Services (SOAP)
 - vex2ccf: Creates control file template from VEX
 - Data providers: Provides chunks of data and transfers them to GridFTP servers.



Implies staged correlation (uncorrelated data hits disk)

Distributed Correlation

Distributed correlation works:

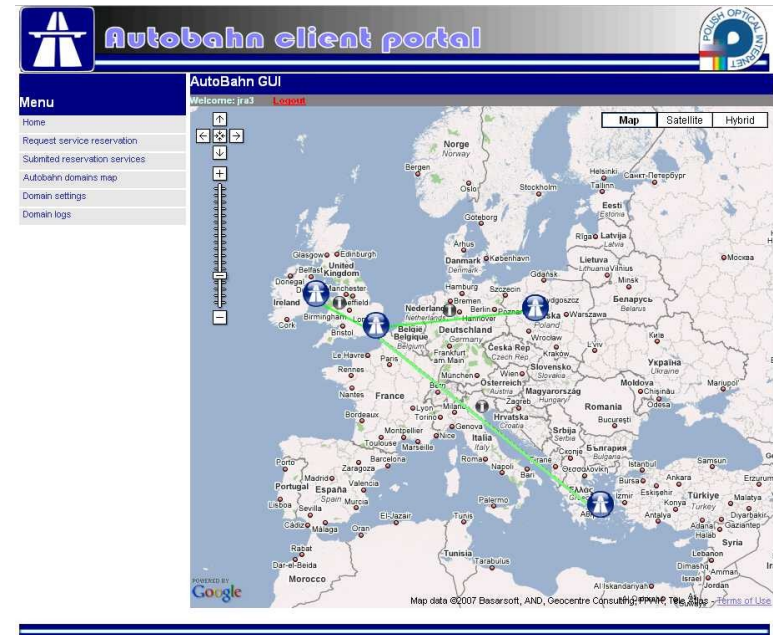
- Successful correlation done on two clusters in Poznan
- Third cluster @JIVE being added

But:

- Grid certificates don't fit our community; use ssh instead.
- Existing HPC clusters are not necessarily suitable:
 - NFS performance problems.
 - Job scheduling seems to favour small jobs.
 - Libraries not always present.
 - Infiniband MPI implementations can be quirky.

AutoBAHN

- JRA within GEANT2 developing a bandwidth-on-demand facility across domains
- Implemented by NRENs of Greece, Ireland, Poland and Croatia, using the GEANT2 testbed
- More NRENs working on deployment (but not SURFNet?)
- Interfaces to other BoD-systems being developed (Internet2).
- Two ways of making reservations:
 - Interactive through web interface
 - Using a webservice




Using AutoBAHN for eVLBI

Interesting for several reasons:

- Many arrays don't observe 24/7
- Correlators in different locations
- Reduce cost?

Successfully demonstrated at SC08:

- Circuits configured directly by correlator
- Data streamed directly into DAS-3 cluster
- Also included data from Boston through Internet2 DCN



Future Development

- Pulsar Binning
- Finish GPU evaluation
- Add VDIF support
- Add jive5a support

Conclusions

- SFXC is ready to do real science.
- AutoBAHN is a promising technology.
- Can't run on just any cluster.
- Offloading FFTs on GPU is not worth it.

Thanks

- The GEANT2 AutoBAHN team
- Dominik Stoklosa, Marcin Okon and everybody else involved in FABRIC at PSNC