FABRIC

A pilot study of distributed correlation

Huib Jan van Langevelde

Ruud Oerlemans

Sergei Pogrebenko

and many other JIVErs…

Aim of the project

- Research the possibility of distributed correlation
  - Using the Grid for getting the CPU cycles
  - Can it be employed for the next generation VLBI correlation?
  - Exercise the advantages of software correlation
    - Using floating point accuracy and special filtering
  - Explore (push) the boundaries of the Grid paradigm
    - "Real time" applications, data transfer limitations
- To lead to a modest size demo
  - With some possible real applications:
    - Monitoring EVN network performance
    - Continuous available eVLBI network with few telescopes
      - Monitoring transient sources
      - Astrometry, possibly of spectral line sources
    - Special correlator modes: spacecraft navigation, pulsar gating
    - Test bed for broadband eVLBI research

Something to try on the roadmap for the next generation correlator, even if you do not believe it is the solution…

SCARI*e* FABRIC

- EC funded project EXPReS (03/2006)
  - To turn eVLBI into an operational system
  - Plus: Joint Research Activity: FABRIC
    - Future Arrays of Broadband Radio-telescopes on Internet Computing
      - One work-package on 4Gb/s data acquisition and transport (Jodrell Bank, Metsahovi, Onsala, Bonn, ASTRON)
      - One work-package on distributed correlation (JIVE, PNSC Poznan)
- Dutch NWO funded project SCARI*e* (10/2006)
    - Software Correlator Architecture Research and Implementation for eVLBI
  - Use Dutch Grid with configurable high connectivity
  - Software correlation with data originating from JIVE
- Complementary projects with matching funding
  - International and national expertise from other partners
    - Poznan Supercomputer centre
    - SARA and University of Amsterdam
  - Total of 9 man year at JIVE, plus some matching from staff
    - plus similar amount at partners

Previous experience on Software correlation

- Builds on previous experience at JIVE
  - regular and automated network performance tests
    - Using Japanese software correlator from NICT
  - Huygens extreme narrow band correlation
    - Home grown superFX with sub-Hz resolution

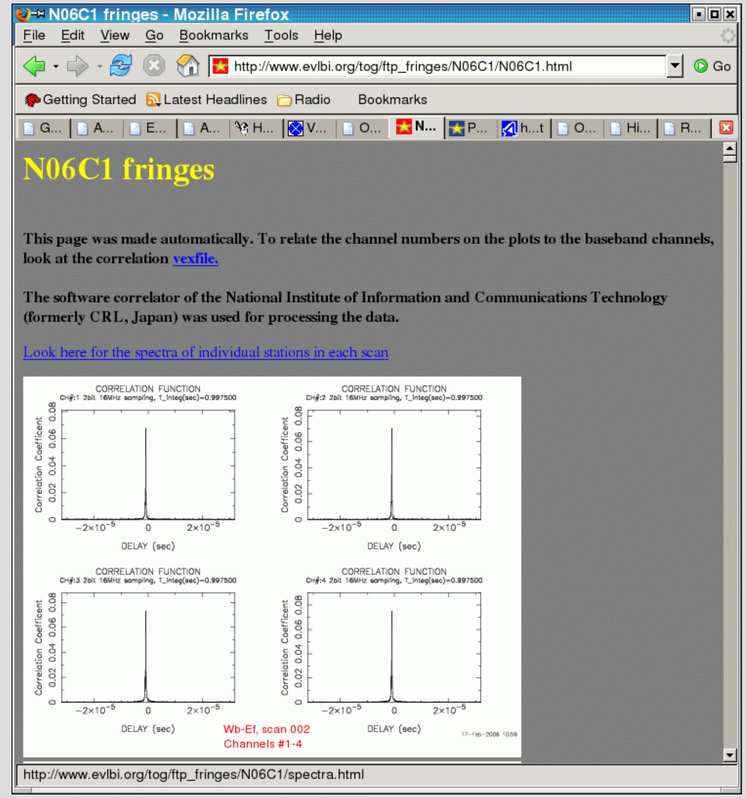

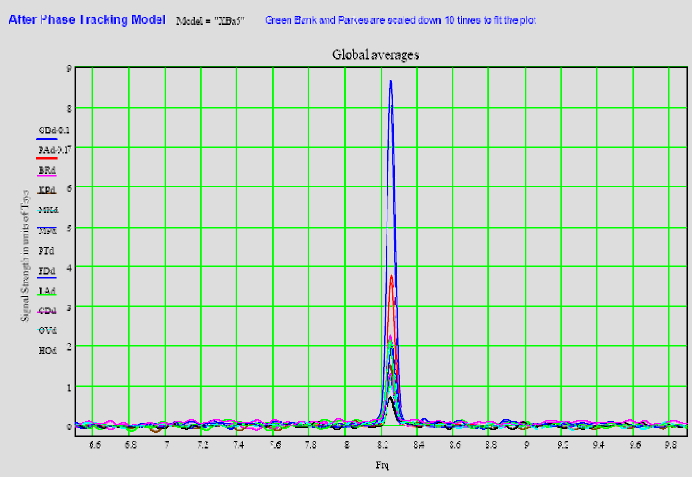

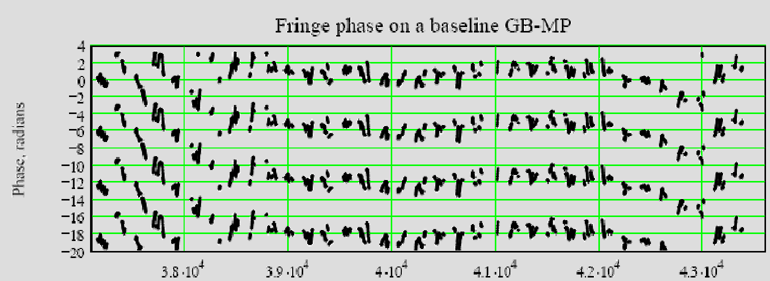

Huygens carrier signal
fringe phases after cross-correlation in 30 mHz wide phase tracking band.

Basic idea

- Use the Grid for correlation
  - CPU cycles on compute nodes
  - The Net could be crossbar switch?
- Correlation will be asynchronous
  - Based on floating point arithmetic
  - Portable code, standard environment

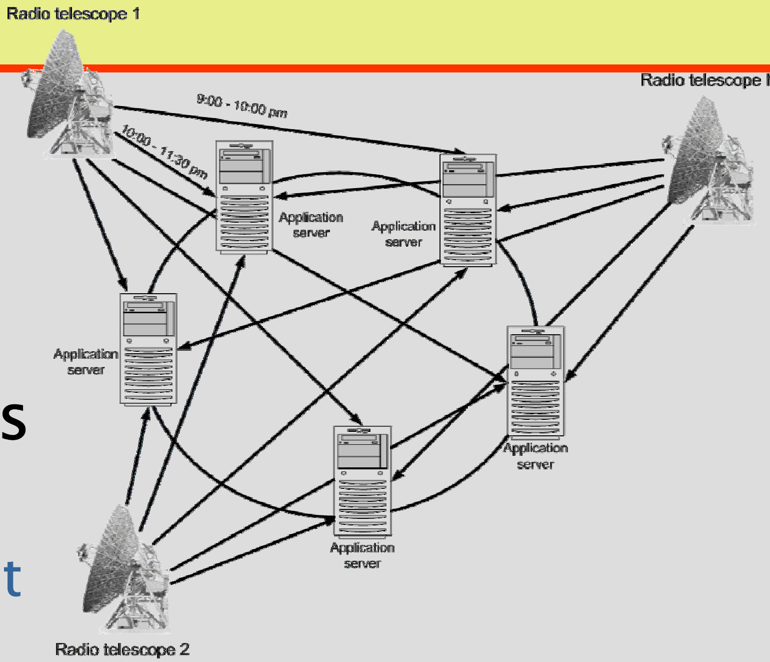

typical VLBI problems					
description	N telescopes	N subbands	data-rate [Mb/s]	N spect/prod	Tflops
1 Gb/s full array	16	16	1024	16	83.89
typical eVLBI continuum	8	8	128	16	2.62
typical spectral line	10	2	16	512	16.38
FABRIC demo	4	2	16	32	0.16
future VLBI	32	32	4096	256	21474.84

Rough estimate based on XF correlation

Work packages

- Grid resource allocation
  - Grid workflow management
    - Tool to allocate correlator resources and schedule correlation
    - Data flow from telescopes to appropriate correlator resources
  - Expertise from the Poznan group in Virtual Laboratories
    - Will this application fit on Grid?
    - As it is very data intensive
    - And time-critical if not real-time
- Software correlation
  - correlator algorithm design
    - High precision correlation on standard computing
    - Scalable to cluster computers
    - Portable for grid computers and interfaced to standard middleware
  - Interactive visualization and output definition
    - Collect & merge data in EVN archive
    - Standard format and proprietary rights

Workflow Management

- Must interact with normal VLBI schedules
  - Divide data, route to compute nodes, setup correlation
  - Dynamic resource allocation, keep up with incoming data!

Effort from Poznan, based on their Virtual Lab.

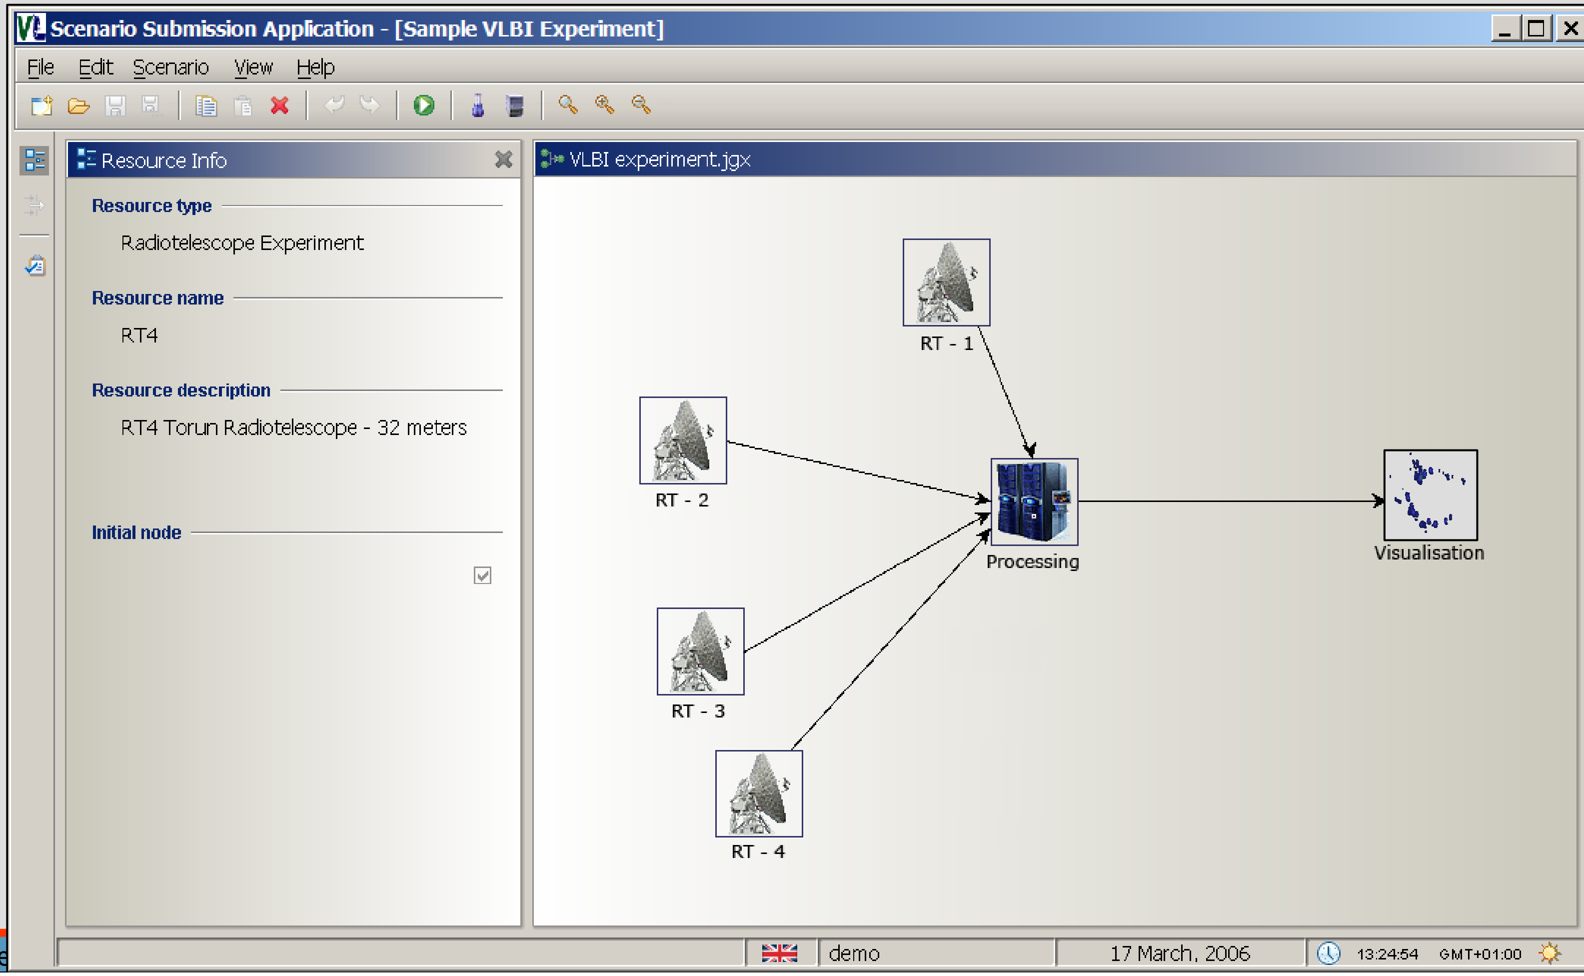

Topology

- Slice in time
  - Every node gets an interval
    - A “new correlator” for every time slice
    - Employ clusters computers at nodes
  - Minimizes total data transport
    - Bottleneck at compute node
    - Probably good connectivity at Grid nodes anyway
  - Scales perfectly
    - Easily estimated how many nodes are needed
    - Works with heterogeneous nodes
  - But leaves sorting to compute nodes
    - Memory access may limit effectiveness

- Slice in baseline
  - Assign a (or a range of) products to a certain node
    - E.g. two data streams meet in some place
  - Transport Bottleneck at sources (telescopes)
    - Maybe curable with multicast transport mechanism which forks at network nodes
    - Some advantage when local nodes at telescopes
  - Does not scale very simply
    - Simple schemes for ½N^2 nodes
    - Need to re-sort output
  - But reduces the compute problem
    - Using the network as the cross-bar switch

Broadband software correlation

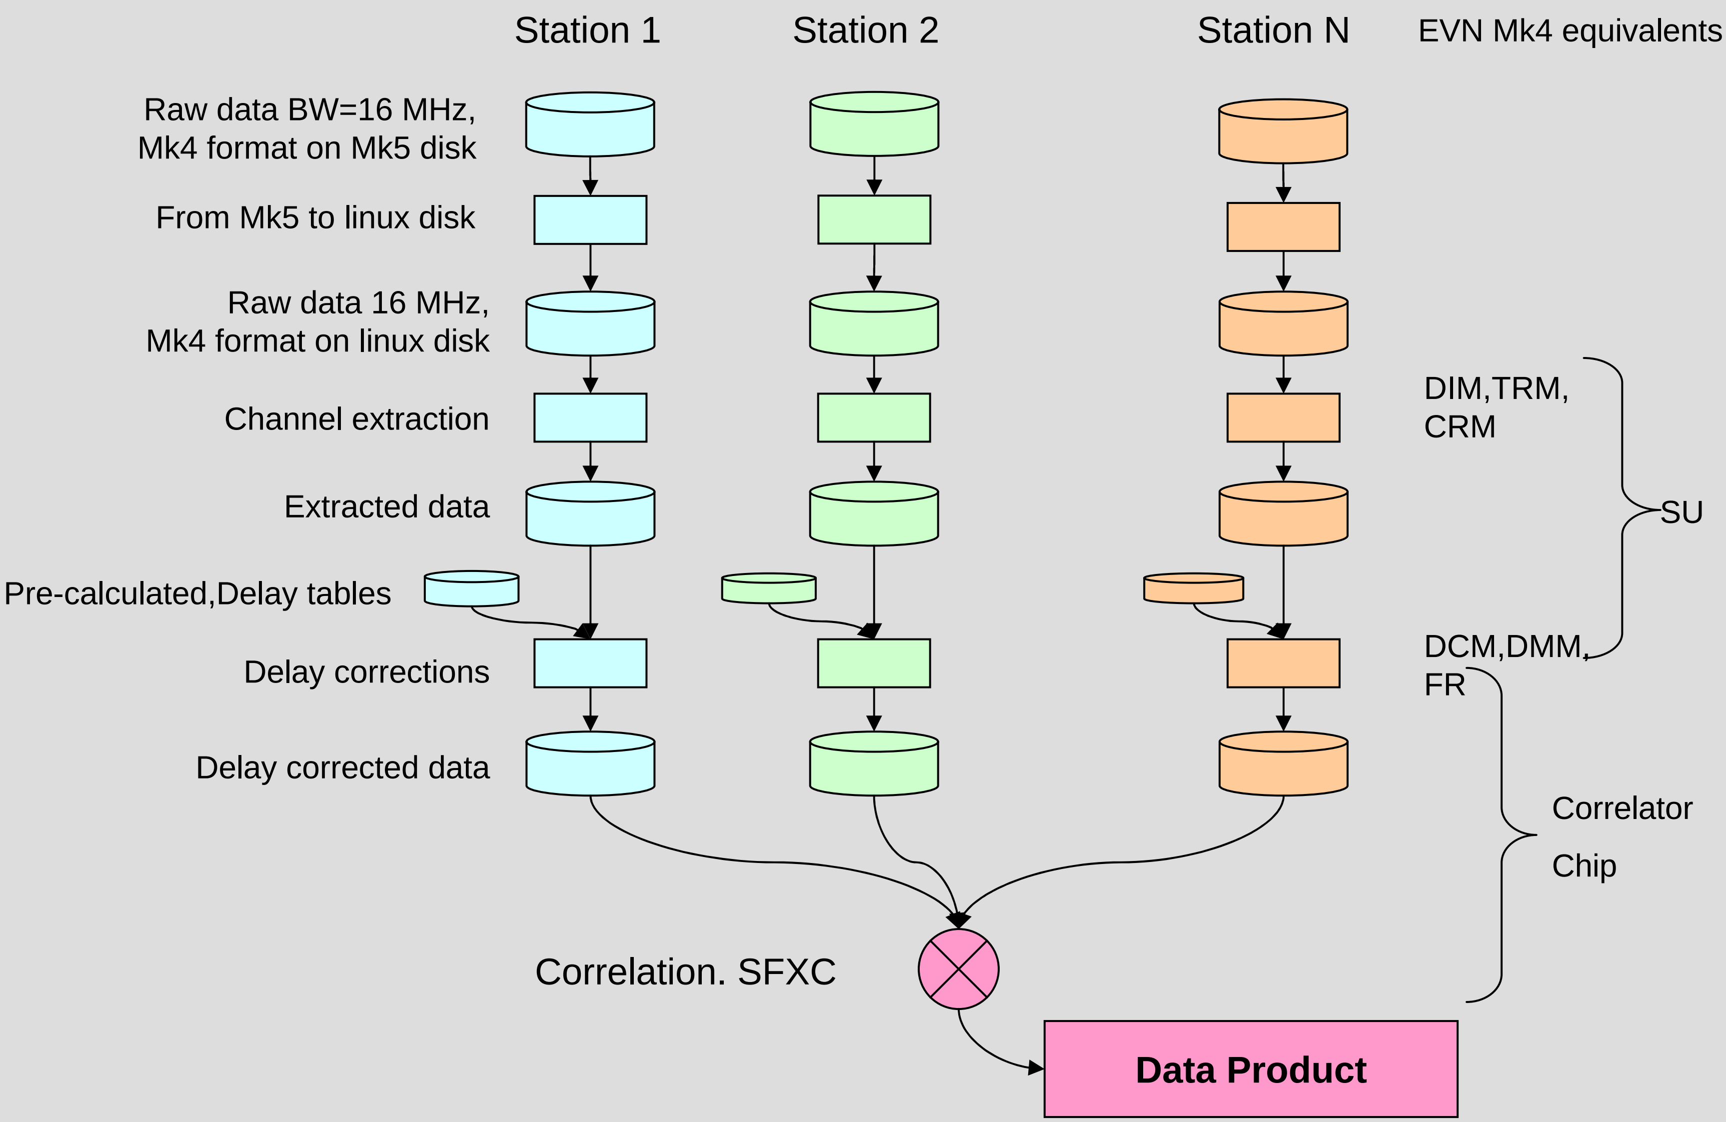

Better SNR than Mk4 hardware

Compare phases of MK4 and Huygens SW correlators, BW 16 MHz, Baseline GB-BR, Source DA193, S-band
210 seconds, 0.5 s integration per point,
64 spectral channels for MK4, 65 spectral channels SW. Bulk linear trend (common slope for both) removed,
10 degrees shift between curves applied for distinction. MK4 data - red, SFX data - blue

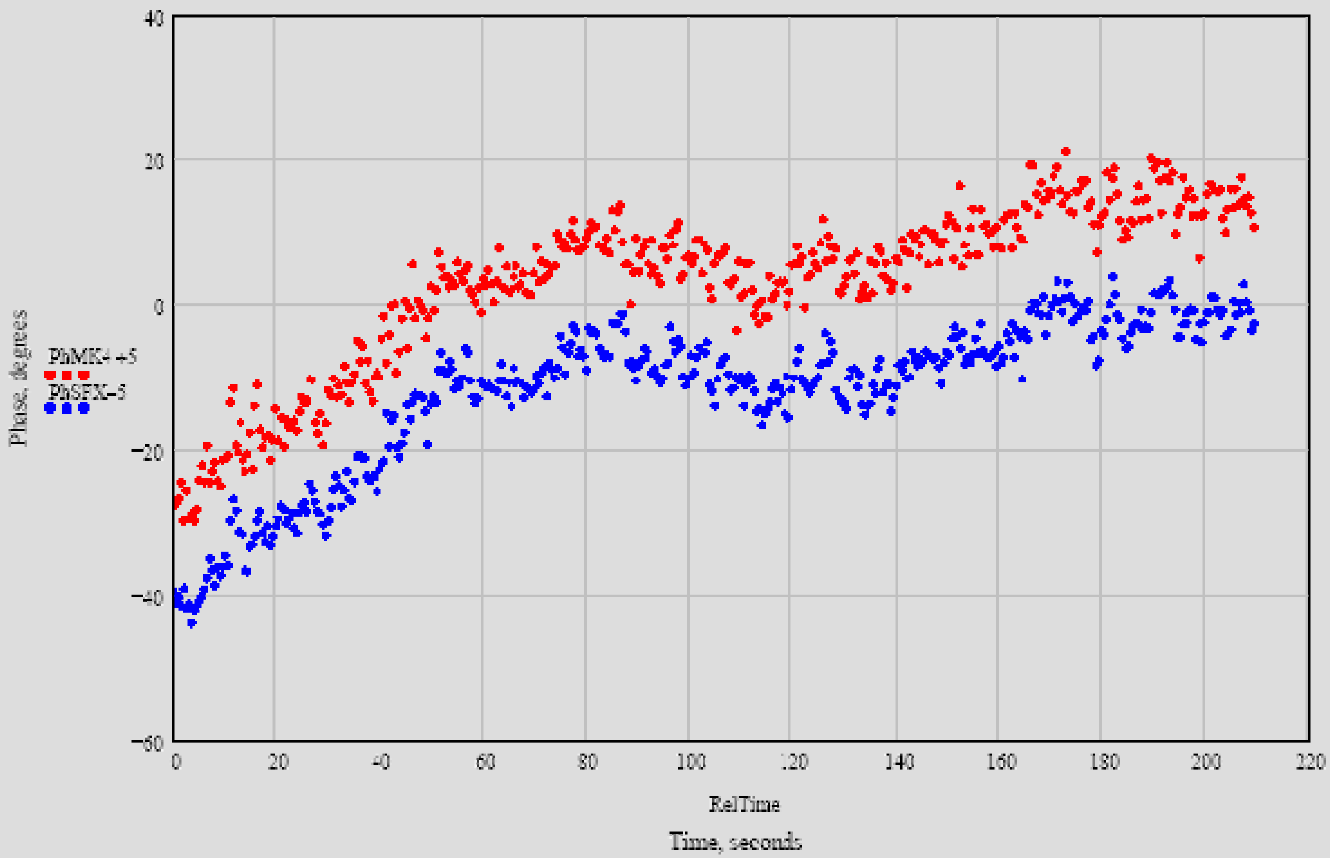

Scatter on SFX data is less than that on MK4

stdev(diffMK4) = 3.1874121493617644

stdev(diffSFX) = 2.2155567957810138

Software correlation

- Working on benchmarking
    - Single core processors so far
    - Different CPU's available
  - Already quite efficient
    - More work on memory performance
- Must deploy on cluster computers
- And then on Grid
- Organize the output to be used for astronomy

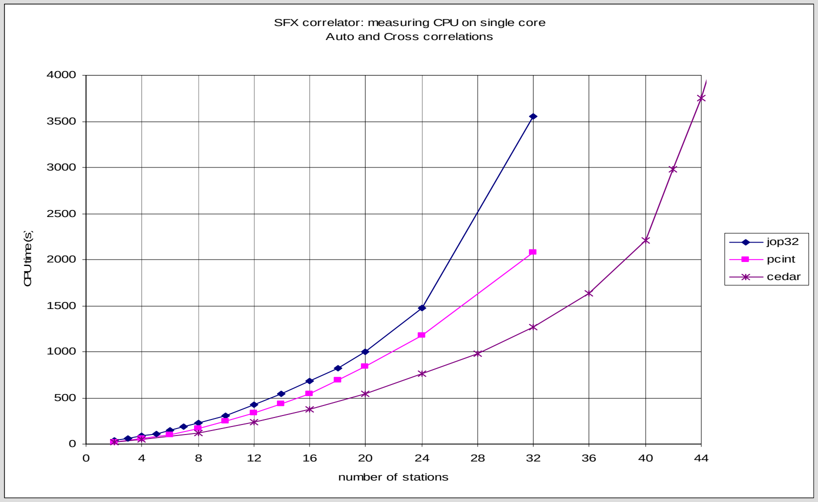

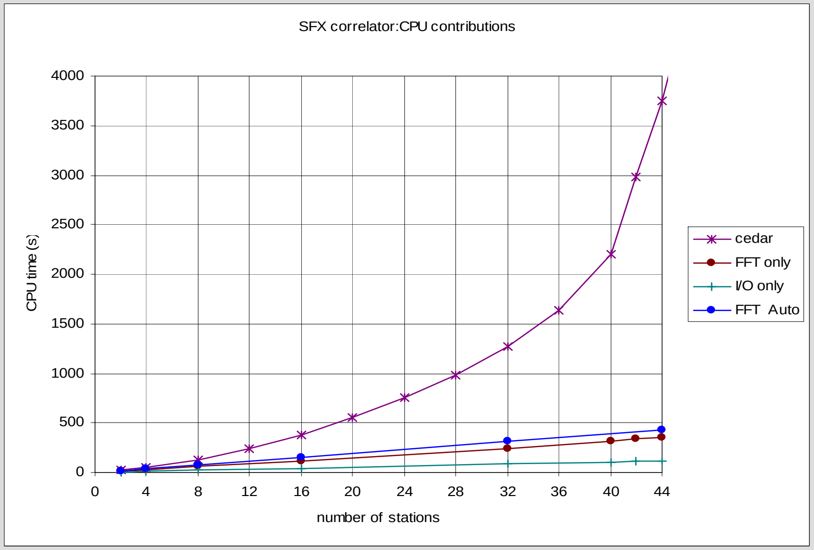

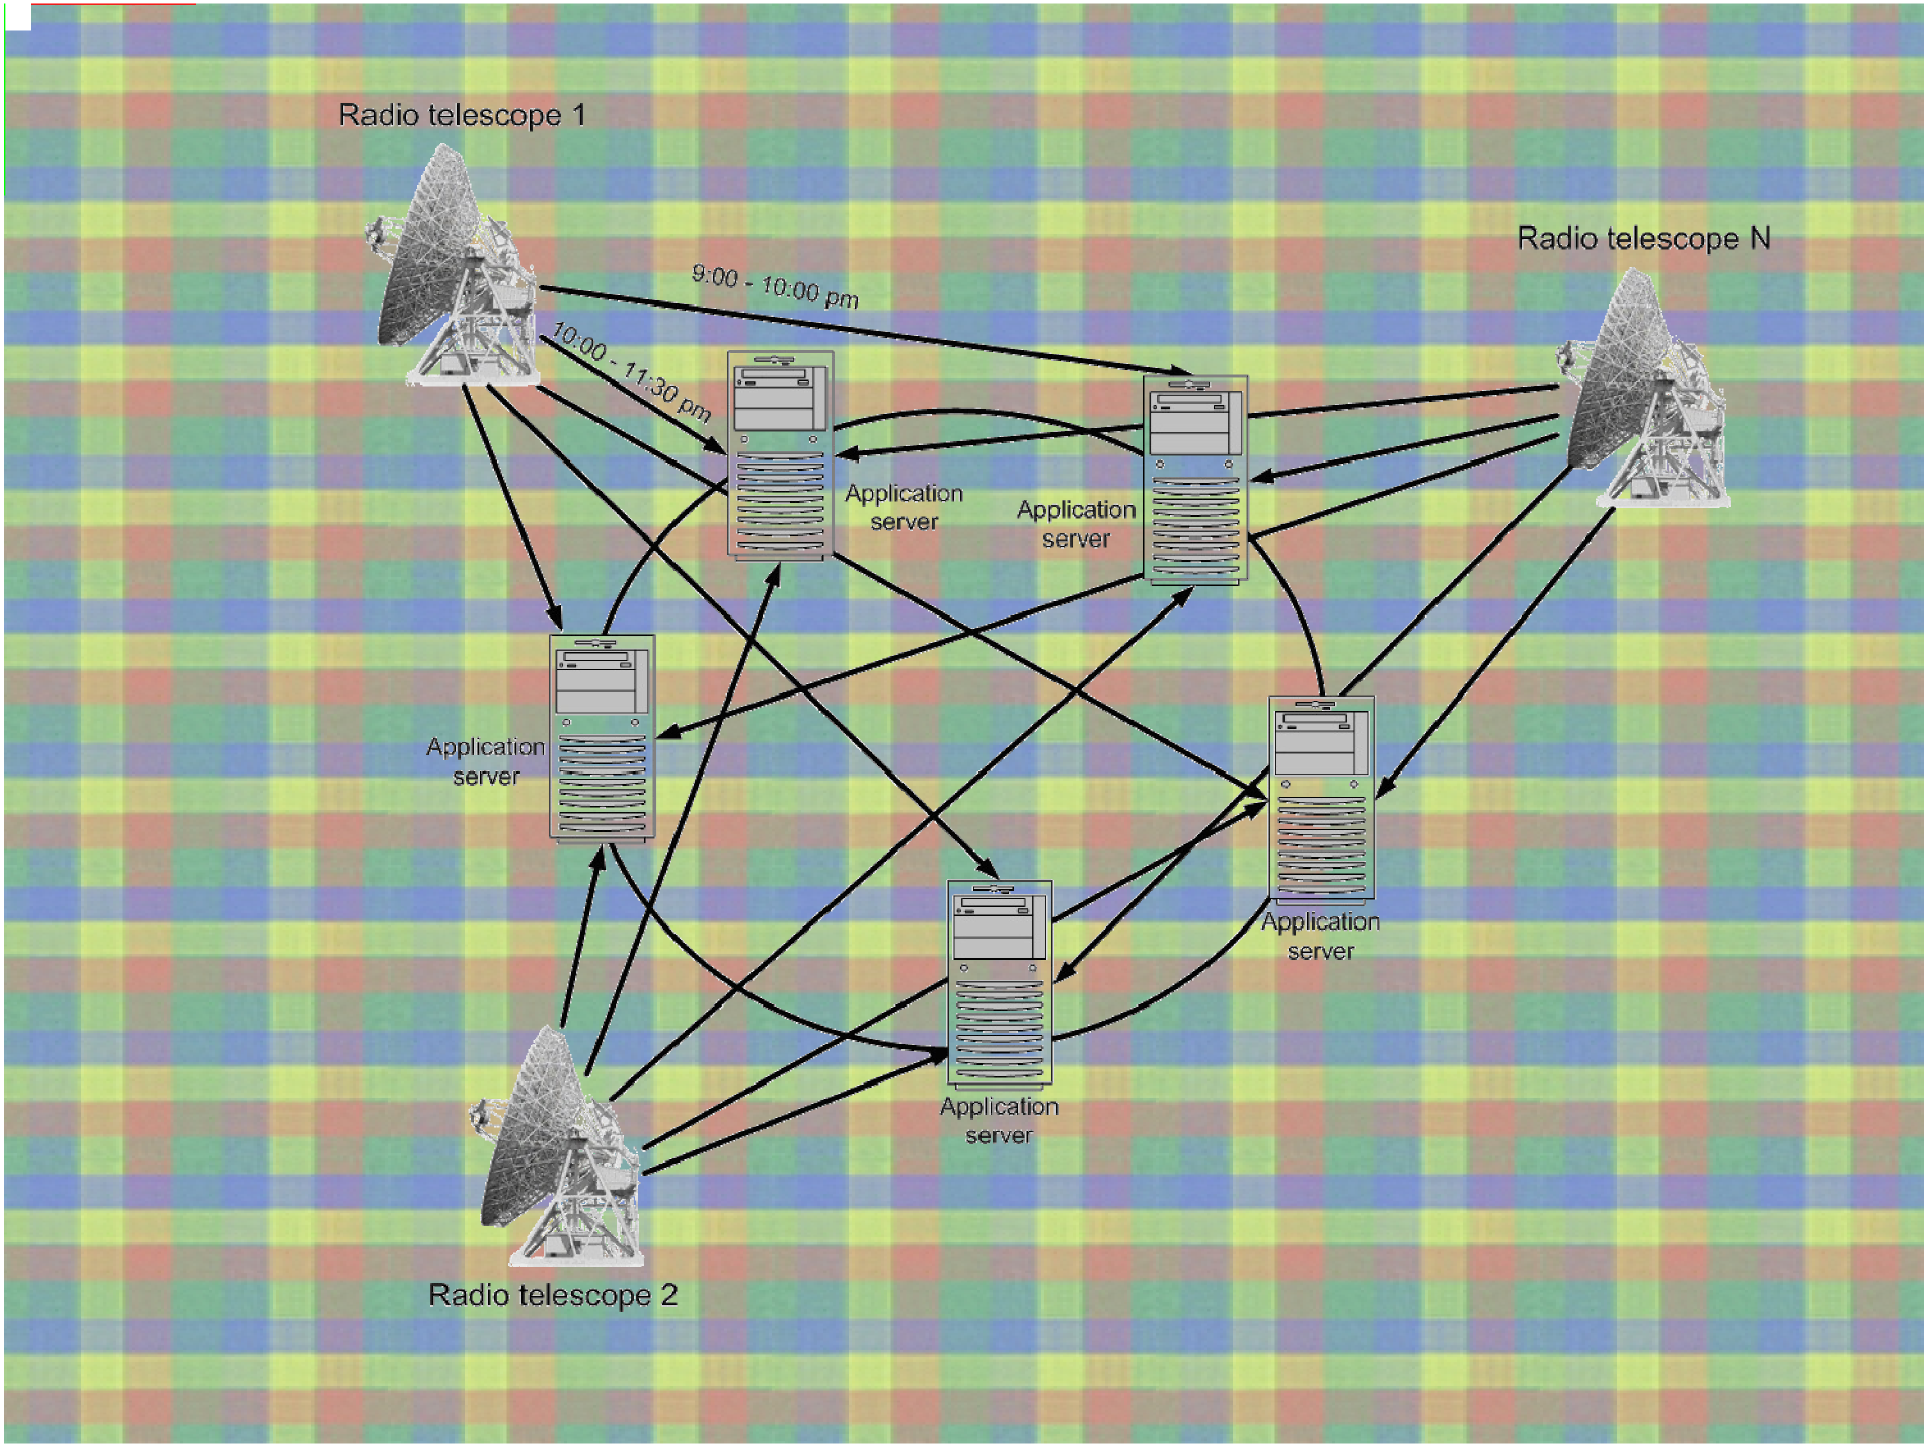

Radio telescope 1
9:00 - 10:00 pm
10:00 - 11:30 pm
Radio telescope N
Application server
Application server
Application server
Application server
Application server
Radio telescope 2