Real-Time System

- CfA
  - Lincoln Greenhill – workpackage manager
  - Daniel Mitchell – calibration & imaging
  - Randall Wayth – imaging & calibration
  - Steve Ord – imaging and survey
- IIC (Harvard)
  - Richard Edgar – GPU coding
- Melbourne
  - Stewart Gleadow – front-end & parallelisation

Outline

- Overview and Status
  - Preparation for a March '09 deployment
- Work to be done
  - GPU conversion
  - Algorithm development
  - Timeline
  - Risk
- Strategy for build-out
  - Porting to a larger cluster
  - Algorithm development
  - Risk

Status of 32T RTS

- All required hardware is in WA.
- Basic pipeline regularly processes MAPS data (CPU version).
- Much of the C-code has been converted to the NVIDIA Compute Unified Device Architecture, CUDA.
- RTS distribution runs in either CPU or CPU+GPU mode (not a full CUDA pipeline yet).
- CPU version has been used to calibrate and image visibilities from the November 2008 site visit. Several hacks need to be worked into the system.

32T RTS Tasks

- Interface between the Correlator and the RTS (Stewart + CfA).
- Head node ⇔ compute nodes code (Stewart + CfA).
- Complete CUDA pipeline:
  - Gridder and imager both implemented (gridder needs optimisation).
  - Calibrator Measurement Loop implemented (needs optimisation and testing).
  - Stokes conversion implemented (needs optimisation and testing).
  - Use simple & fast regridder (bilinear – possibly on the CPU).
- Implement fixes for real data:
  - RFI detection / flagging (currently setting inv_var weights to zero).
  - Direction-intependent calibration using peel solutions of a single bright source.
  - Variable calibrators, resolved calibrators, ...
- Much testing once real data in flowing!

32T RTS Timeline

- January:
  - Finish master node.
  - Finish conversion to CUDA (expect speed up factor of 5-10).
- February:
  - Attach to test-bench at Haystack (integration of UMelb & CfA code).
  - Continue to optimise routines.
  - Continue to test and develop real-time calibration.
- March:
  - Install code on the machines in Boolardy.
  - Test, fix, test, fix, ...

32T RTS Risk ≈ GPUs

- I/O Concerns:
  - INPUT ✓ OUTPUT ✓ (sustained transfers ?)
  - CALIBRATION ? ✓
- Memory Concerns ✓
- Processing Concerns:
  - 10 calibrators → < 65 Gflops/8sec/GPU = 57 Gflops/sec/GPU (7 cycles/flop). ✓
  - Mostly Stokes conversion (fewer baselines than 512T but similar image size).
  - Much of the Stokes conversion code can be pre-computed every few minutes.
- Ionospheric Faraday rotation ?
- Primary beam correction ?

Tasks (build out from 32T)

- Continue to optimise algorithms.
- Primary beam models
  - Stability of beam shapes.
  - Time and frequency derivatives.
  - Parameterisation.
- Ionospheric models
  - Smallest and largest scales: temporal / angular / refractive / ...
- Incorporate Justin Kasper's ionospheric FR technique.
- Investigate various gridding kernels for different science objectives (best case / worst case effect on image quality, deconvolution complexity, ...)

Risks 1 (build out from 32T)

- GPU performance less than anticipated
  - Will learn from 32T.
  - Reduce number of calibrators, reduce baselines used in the fits, ...
  - Simplify regridding / Stokes conversion.
- Not enough calibrators
  - **Limited FLOPS**: increase FLOPS.
  - **Limited by confusion noise**: use a more complicated model for the pre-peel step.
  - **Small scale structure in the ionosphere**: use sources from the images.
  - Key science applications require 100s of sources subtracted in real-time.

Risks 2 (build out from 32T)

- Variable calibrators
  - **Spectral variations**: 5-10% errors from catalogues, or determine during absolute calibration process.
  - **Temporal variations**: model variation, or don't use in primary beam fits.
  - **Ionospheric scintillation**: detect and don't use in primary beam fits.
- Other real-time applications
  - Need to manage real-time resources.