Casey Law
RSRO final report
23 April 2012

Summary:
The central goal of my RSRO was to decrease the WIDAR integration time to much less than the standard of 1 second. The primary science goal is to use correlated data products to detect sources like pulsars on 10 ms time scales. This time scale also corresponds to the dispersive delay at gigahertz frequencies. Thus, writing 10 ms integrations makes it possible to detect fast transients and measure dispersion with correlated data for the first time with the JVLA.

Generally, this goal focused on measuring the data rate at which WIDAR can write fast dumps. This also has applications modes for slow time scale, spectrally-large data. We ran tests throughout my visit to test the rate under a range of correlator configurations. These tests revealed a failure modes that were: related to the network connection from baseline boards to the CBE, within the CBE software, and in the writing of data to the Lustre file system. These tests also helped us identify an underlying theory for each bottleneck. This is important for understanding how it affects other modes (e.g., large spectra) and how future upgrades can improve the performance of WIDAR.

I assisted with tests run by Michael Rupen, Ken Sowinski, and Martin Pokorny. Early tests revealed that the correlator could write data at a rate of 25 MB/s, corresponding to 100 ms integrations in the full array with 2 subbands and 2 polarizations. Pushing the system to shorter integration times revealed multiple failure modes that led us to conclude that the CBE software was the strictest bottleneck. Martin then optimized the correlator software to reduce the number of processes that may block each other. This improved performance dramatically. The current fastest observing mode writes data at a rate of 110 MB/s, which corresponds to 24 antennas at 10 ms rates, 2 subbands, and 4 polarizations. Other tests suggest that we can easily expand on this limit.

This is the first step in a long term goal to enable real-time transient detection with the JVLA. The CBE is a flexible environment that has access to all the visibility data needed for this goal. However, since the RSRO is so central, it does a lot of work and is home to many critical bottlenecks. By increasing the data rate written by WIDAR, we have identified how much room there is for other kinds of computation on the CBE.

Products and results of the RSRO:

1) Ability to write 10 ms visibilities with the full array -- We have shown that we can do this with 2 polarizations and 3 subbands with the full array.

2) Observe with fast dumps as part of RRAT science program -- We used the mode described above to observe one RRAT for 30 minutes; the full program will observe for 15 hours. For now, this needs to be scheduled and run manually by Michael or Ken.

3) Describe WIDAR tests and bottlenecks -- Our tests identified the main bottlenecks for high data rate configurations. These tests will be summarized in a spreadsheet and in a new WIDAR bottleneck diagram. These products will be created in collaboration with NRAO staff. (Links to this stuff coming soon.)

4) Identify potential upgrade path -- The correlator speed test results show how the NRAO can most efficiently invest in the correlator to improve its performance. In general, the CBE is a very powerful and central component of the system. Software changes have improved performance dramatically by improving parallelism. This means that expanding on the number of nodes in the cluster will likely improve performance further, though not necessarily in proportion to the number of nodes.

Work in progress:

1) Normalizing fast-dump observing -- The fast dump mode still needs to be manually configured and started. Discussions with Sonja and others from the HIA are ongoing to help simplify this process.

2) Build software for CBE for real-time, commensal transient detection -- This is part of the long term goal for this project and remains unfinished (and unfunded).