

# Softc: an Operational Software Correlator

*Stephen T. Lowe*

*Jet Propulsion Laboratory / California Institute of Technology*

*e-mail: [steve.lowe@jpl.nasa.gov](mailto:steve.lowe@jpl.nasa.gov)*

## Abstract

Softc has been used operationally for spacecraft navigation at JPL for over 2 years and will be JPL's Mark 5 correlator next year. Softc was written to be as close to an ideal correlator as possible, making approximations only below  $10^{-13}$  seconds. The program can correlate real USB, real LSB, or complex I/Q data sampled with 1, 2, 4, or 8-bit resolution, and was developed with strong debugging tools that made final debugging relatively quick. Softc's algorithms and program structure are fully documented. Timing tests on a recent Intel CPU show Softc processes 8 lags of 1-bit sampled data at 10 MSamples/sec, independent of sample rate.

## 1. Introduction

Level-one VLBI processing has traditionally required state-of-the-art processing power and data bandwidths that could only be implemented using custom electronic hardware and data storage equipment. These devices, correlators, often took many years to design and build followed by at least a year of use to excise the more obvious processing problems. This long development time also ensured that by the time a correlator was debugged and in a stable operational mode its equipment was obsolete. Because there was no other processing option, the VLBI community has endured high costs associated with these devices. These costs include many millions of dollars for development, on the order of \$1M annually to pay for the ongoing debugging, capability improvements, and maintenance, a trained operator staff, and the infrastructure costs of the correlator room such as leasing, power, and air conditioning. Data storage and transport has used custom tape technology, but this is also expensive, and as the technology has been pushed to higher densities, tape drive/head maintenance costs have increased. Finally, some indirect correlator costs include lack of flexibility for novel experimental setups, lack of transparent processing algorithms, lack of repeatable output due to tape playback errors, and downtime for maintenance, testing, and modifications.

It has been known for some time that all the disadvantages of the traditional correlator noted above could be eliminated if efficient data handling could be moved from tape to disk, and general-purpose computers could meet the processing requirements. Until recently, neither of these conditions could be met adequately, but interestingly it now appears the crossover points for both of these competing technologies, custom tape vs. disk, and custom hardware vs. general-purpose computer, is in the recent past [1]. Both the recent Mark 5 disk format and software correlators such as Softc take advantage of huge commercial R&D budgets worldwide. These are cases of custom VLBI hardware capabilities being overtaken by commercial products developed with enormous resources compared to those in the VLBI community. By extrapolation we are likely to experience other paradigm shifts soon in the remaining custom VLBI hardware, namely station electronics. We are also beginning to see how all these changes are significantly altering the optimum values

of VLBI system parameters. For example, it will be more efficient to have smaller antennas with higher bandwidths, and reliability should increase with the improved hardware, disk storage, and real-time fringe verification.

## 2. History

In 1995 JPL began to decommission its Delta-Differenced One-way Range (DeltaDOR) spacecraft navigation system. The decommissioning process was to mothball the DeltaDOR capability, eliminate its near real-time Block I hardware correlator, and after a short time cease DeltaDOR funding. At around this same time, test programs were written to assess software correlation speed and to find fringes in old Block I quasar data. These tests showed the Block I bandwidths were low enough that it could be replaced with a software correlator. Since there was no other choice to preserve this capability, a project to replace the Block I with a software correlator began in 1996. This task was essentially completed but since this was a mothballing effort and DeltaDOR was not necessary for any mission, the program, called Softc, was never used. This was probably fortunate as the program, by design, exactly duplicated the Block I processing including its many approximations and flaws. In 1998, the RadioAstron project agreed to fund further development of Softc as a debugging tool for their project. This opportunity was used to essentially start over and build a true software correlator with all its intrinsic advantages without being constrained to a hardware correlator's shortcomings. Unfortunately, this effort was canceled in 1999 just as full code testing began. Softc then remained complete but unused and untested until 2001 when the mis-navigation and demise of Mars Climate Orbiter prompted JPL to resurrect the DeltaDOR navigation technique. Softc underwent full-code testing using geodetic experiments and DeltaDOR passes of Mars Global Surveyor, in orbit at the time, and Mars Odyssey, which was on its way to Mars at the time. Softc's first critical use was the successful Mars Odyssey orbit insertion in Oct 2001. Softc has also been used to successfully navigate both Mars Exploration Rover missions, Deep Space 1, Europe's Mars Express, and Japan's Nozomi and Muses-C (Hayabusa) missions.

## 3. Capabilities

Softc was designed to be as close as possible to an ideal correlator, where ideal refers to processing accuracy; no compromise in accuracy greater than  $10^{-13}$  sec was made to increase performance (or for any other reason). Softc can correlate 1, 2, 4, and 8-bit sampled data, upper, lower or double (I/Q) sideband data, and data using either of two sample encoding schemes. Softc is quite portable, as it was developed on a little-endian Alpha running VMS, it works operationally on a large endian machine running Solaris, and has been tested on little-endian Intel machines running Linux. It produces identical output on these machines with no code modifications, even with different C compilers.

Softc was designed to process essentially any VLBI data. To do this and remain independent of the hardware and post-correlation software interfaces, Softc has its own input and output data formats. These formats were designed to be as simple as possible, but a translation program is needed to convert the sampled data into Softc's format, and another program is needed to translate from Softc's output to the desired post-correlation format. It is also possible to place these translators inside Softc; this has been done on the output side for our DeltaDOR effort and will likely be done for Softc's input in the future.

## 4. Code Debugging

One of the greatest hurdles in developing a VLBI correlator is the elimination of processing errors and inaccuracies. It is common for hardware correlator developers to spend years tracking down reasonably significant bugs. Also common are correlator users finding new, consequential errors long after all known problems were eliminated. An important limitation for correlator developers has been the lack of good debugging tools. For this reason, a significant effort, perhaps 40% of Softc's development, went into the creation of a full Monte Carlo data generator. This generator creates simulated VLBI data that can then be processed by Softc. In this way, the expected results can be calculated exactly a priori, and any deviations from these results indicate problems with the correlator.

The Softc Monte Carlo can simulate data in two ways. The first is an engineering mode where data can be generated with user selected delays, delay rates, fractional sample offsets, etc. This mode was very useful in the initial stages of debugging, as the implied geometric models were quite simple and covered a much greater parameter space than could be achieved by Earth-based station models. For example, data can be generated having a constant geometric fringe rate many times that possible for any Earth-based experiment; Softc's performance under these extreme conditions is a good check that fringe-frequency related calculations are done correctly.

The second Monte Carlo mode simulates data using the correlator model. Correlating this simulated data with Softc using the same model should result in no residual delays or phases, and tests Softc in its usual processing mode. The power of this type of test was confirmed when data were simulated using the Block I correlator model for an old experiment. When the Block I processed the simulated data, the residual delay was not zero, but a constant 1-sample delay. In other words, it appears the Block I had always reported delays with a constant 1-sample delay error, and this problem was never found throughout the life of the correlator, but was found immediately with this Monte Carlo test. Softc has passed a large number of extensive Monte Carlo tests, and this, combined with the fact that few calculational approximations are made, gives great confidence that any remaining processing errors are either insignificantly small or easily corrected.

Other tools were also created to find errors early in Softc's development. Because Softc was originally created to correlate 1 and 2-bit sampled data, the core processing routines perform extensive bit-level manipulations. Several test driver programs were developed for each bit-manipulation routine to check its function. Although these drivers are not part of the run-time portion of Softc, they should be considered part of the code package in the same way the Monte Carlo is part of Softc. These programs were designed to test the bit-manipulation routines under the most extreme conditions, and a significant effort went into trying to find processing errors in the core routines. The results of these tests were that the core routines, which perform the bulk of the processing, could be considered essentially bug-free. This made full-code debugging much easier because when a strange problem was seen in the results, many of the most perverse explanations one might imagine could be dropped from consideration.

In full code testing we found the unresolved errors to be dominated by modeling errors. The Monte Carlo cannot detect such problems because the same bad model both generates and correlates the data, or in other words, Softc cannot assess what is or is not a realistic model. For example, the first significant problem found turned out to be a sign error in the troposphere model. The Monte Carlo did not detect this (and could not possibly have done so), so only full code testing of real data could find these types of errors.

## 5. Future Directions at JPL

We recently obtained a new 18-CPU Beowulf computer and are currently porting Softc, Fit (our post-correlation software), and Modest (our parameter estimation software) onto this machine. We plan to interface this to two Mark 5 recorders by Jan 2004 and this will be JPL's VLBI correlator. Work is in progress to update Fit and to add additional Softc capabilities so we should have a clean, modern software system in early 2004. Timing tests on this system show Softc can process 8 lags of 10 MHz sampled real USB single-bit data in real time, and that the processing time is roughly linear with sampling rates from 1 to 100 MHz.

We are canceling the development of our 4-station hardware correlator, which is now essentially complete. There is nothing wrong with this correlator other than it is a *hardware* correlator with all its intrinsic disadvantages. Walking away from our hardware correlator now puts us on a path with a future and will likely save significant money. Because the economics of this decision are probably common to other institutions, we may not see another 4-station hardware correlator being developed anywhere. The economics may also favor replacing larger correlators even now, but the community probably needs to go up the software correlator learning curve beginning with the smaller systems.

Work is under way to obtain an open software license for Softc so that others can benefit from this program. Softc and its algorithms are documented [2] and that report should be publicly available very soon.

## 6. Summary

Softc was created to be as accurate as possible, capable of processing essentially any VLBI data, pass strong debugging tests, have a simple user interface, have no platform dependencies, and be written modularly in the most common programming language. It has been used operationally for spacecraft navigation for over 2 years and will be JPL's Mark 5 correlator next year.

## References

- [1] See Whitney, R., "Mark 5 Disk-Based Gbps VLBI Data System," MIT Haystack Observatory, <http://web.haystack.mit.edu/mark5/software.html> for more on the tape vs. disk issue.
- [2] Lowe, S., "Softc: a Software VLBI Correlator," JPL Report, in press.