A VLA Experiment – Planning for Voyager at Neptune

J. W. Layland
TDA Planning

P. J. Napier and A. R. Thompson
National Radio Astronomy Observatory

A VLA engineering experiment was conducted on the night of July 22, 1983 to explore one aspect of the potential for the VLA to support Voyager at its Neptune encounter in August of 1989. Specifically, the experiment tested the ability of the VLA to self-calibrate on a natural radio source whose effective signal strength is the same as Voyager's will be at its Neptune encounter. The experiment was successful and supported the belief that the VLA would be able to be self-calibrated with Voyager's signal.

I. Introduction

The Very Large Array (VLA) is an array of twenty-seven 25-meter antennas in a triradial configuration in the high New Mexico desert. Reference 1 provides a comprehensive overview of the VLA. Details on the VLA, its hardware and software, operation, and the theory behind it may be found in Ref. 2. The primary role of this array is developing maps of radio-bright objects in the sky, and it incorporates a large mapping processor which is capable of cross-correlating the 351 (= 27*26/2) baselines of the array in real-time. One of the optional products of this mapping processor is a combined output which represents the coherent sum of the signals being received at each of the antennas. This combined output of the mapping processor would represent about 2-1/2 of the DSN's 64-meter apertures if all of the VLA antennas were outfitted with X-band cryogenic field effect transistors (FET's). As noted in the Interagency Array Study Report (Ref. 3), this capability can be of significant benefit to Voyager at its encounter with Neptune in 1989.

Configuring the VLA to be able to receive signals from Voyager will require significant investments of time and money. One of the many questions asked before seeking to use the VLA was whether the VLA could self-calibrate on the signal from Voyager, and thus be able to easily combine the signals from all 27 antennas without recourse to a stronger nearby natural radio source as a calibrator. As the VLA could not receive at X-band, obtaining a preliminary answer to this question entailed observing some radio source which provided a signal in one of the existing VLA receive bands which was a suitable analogy to the signal expected from Voyager at Neptune in 1989.

Voyager's total signal from Neptune is predicted to be 3.6×10^{-21} W/m^2 (B. D. Madsen, Voyager Telecom, private communication, May 2, 1983). If we filter it into a 60 kHz bandwidth, selecting one sideband and maximizing the bandpass SNR, the signal is cut by 5 dB to 1.1×10^{-21} W/m^2. Since Voyager's signal is matched to the feed polarization and will appear in only one of the two polarizations (right-hand circular, but not left-hand circular), while the total flux of a natural radio source is measured in both polarizations, the total cataloged flux of the comparable natural radio source is 3.6 Jy. The VLA processor develops integrated cross products each 10 seconds (or a multiple thereof), which are used to estimate the differential phases between antennas. These cross products do not differ in character whether the input bandwidth is 60 kHz or 50 MHz. A natural source of 130 mJy strength in the full VLA bandwidth of 50 MHz should provide the same SNR in the integrated cross-products as a 3.6 Jy source in 60 kHz and, hence, the same as Voyager's signal. This is significantly below the strength at which the VLA is normally calibrated.

In the experiment itself, radio sources at both 0.5 Jy and 0.14 Jy were included in the observation catalog to help expose any sensitivity to source strength at this level. The observing control "deck" was set up using the time sharing PDP-10, and then called by the array operator into the Modcomp real-time system when needed that evening.

The observation interval occupied 3 hours and contained the observations listed in Table 1. All observations were taken in the 6-cm band, with a nominal system temperature of 60 K. Except where noted, the integration time is 10 seconds and the closed-loop tracking of the differential phases was performed with a loop gain of 1/4. All observations appeared to lock-in solidly, and suggest that self-calibration on still lower source strengths would be quite possible. The estimated performance of the phase tracking operation was extracted by post-processing the recorded data in the VLA computer.

II. Experiment Description

For most of the experiment the VLA was operated with its fundamental integration interval set at ten seconds. Each ten seconds, the 351 complex integrated cross products of the signals received at each of the 27 antennas were recorded for final processing. Also each ten seconds, the 26 cross products for each of the antennas paired with the selected reference antenna were sampled and used to estimate a differential phase which can represent atmospheric shifts, uncompensated equipment drifts, source position offset, etc. When the array is being calibrated, this differential phase is used as the error signal input for a very narrow-band first-order sampled-data phase-locked loop. In normal operation with a 10-sec integration time, the loop gain is set to 1/4, which results in an effective memory time of about 40 sec, or a bandwidth of 0.025 Hz. For the purposes of the engineering experiment being described, the final processing of the recorded data is devoted to estimating how well this sampled-data phase-locked loop was able to track the actual phase variations.

The final non-real-time processing of the data is performed by the VLA program *ANTSOL* (Ref. 2) which uses the stored integrated cross products to perform a global least squares estimate of the apparent gain and phase of each of the 27 antennas. This is the relatively well-known multi-parameter form of the simple least-squares estimator, which will be used but not developed here. The program uses all 351 of the antenna pairs to provide a global solution for all 26 of the phases measured relative to the reference antenna, which is better by a factor of 2/27 smaller variance than the phase estimates based upon a single integrated cross-correlation. The global solution can be obtained using the elementary (e.g., 10-seconds) integration intervals which were sampled in the real-time processing, or the elementary integrals can be accumulated to increase their SNR before the global solution is done. Assuming there is no significant external perturbation of the relative phases, their sample variance is a measure of the performance of the real-time phase tracking process.

All of the data from the July 23rd engineering experiment were non-real-time processed with an integration time of 120 seconds. A small part of the data was also processed at 30 seconds to improve visibility into the phase-estimator behavior. The approximate relation between the observed phase variance in the final processing, and the variance of the phase attributable to real-time processing is calculated in the Appendix.

III. Observations

As listed in Table 1, the experiment consisted of 10 observations of 6 distinct sources, with distinguishing characteristics to be noted in the following. Sources are identified by their standard cataloged positions, with the first source, 1803+784 being at right ascension of 18 hrs, 03 min, and at declination +78.4 degrees. This source is one of the VLA calibrators, and its recorded flux at 2.52 Jy was used as the reference against which the strengths of other, weaker sources were compared.

The time of this first observation was 04:12 to 04:21 (IAT), and the elevation angle at which it was observed was approximately 45 degrees. Final-processing of this source with a 2-min integration time resulted in an observed phase jitter of 2.5° rms. Converting this into an estimate of the jitter in the real-time control loop by the methods of the Appendix yields a value of 3.7° rms for the phase-tracking process with a 10-sec T_I and loop gain of 1/4. The data from this source was not processed with 30-sec intervals.

For each source, the very beginning of the recorded data is ignored until the acquisition transients have subsided. Then, as many 120-sec (30-sec) segments as will fit into the allotted time are accumulated and the phases and effective gains are estimated for each of the 27 antennas, using the 351 cross correlations between antenna pairs. The phase parameter of the selected "reference" antenna is arbitrarily fixed at zero. For each segment, the rms of the 26 non-reference phases was calculated. The average, and one-sigma, uncertainty of this rms value was then calculated over the (10 to 20) segments and appears in Table 1 in the columns identified as RMS PHASE. The column identified as SIG-PH is the estimated rms phase variation in the real-time phase tracking process as derived from 120-sec segments by the methods of the preceding section.

The second, third, and fourth observations were all processed with both 120-sec and 30-sec segments. The rms phase at 30-sec integration time is approximately the same as the rms jitter in the real-time control loop. In each case this estimate is larger than that developed on the basis of the 120-sec segments.

These three observations (numbers 2 through 4) were planned to compare the behavior at high (78°) versus low (16°) elevation, and at about 0.5 versus 0.14 Jansky source strengths. It is clear from this set that phase jitter is much more dependent upon elevation angle than upon source strength – at least for the levels at which we need to be concerned. There are presumably variations of the atmospheric path lengths that dominate the observed phase jitter.

Later observations number 7, 8, and 10 used the same 0.14-Jy radio source with varying real-time processing strategies. For number 7, the unit integration time was doubled, to 20 sec, which resulted in an apparent 20% decrease in phase variation, as compared to a 40% decrease which could theoretically have resulted if the source of the phase variation was independent between the 10-sec segments. Further, lengthening the segment time to 60 sec (and correspondingly increasing the loop gain to 1/2) actually resulted in an apparent 10% increase in phase variations. The expected response if the sources of the phase variations were independent between the 10-sec segments would have been about a 25% decrease.

While these changes are not significant in a statistical sense because of the wide uncertainty of the observed phase variations, they strongly suggest that lengthening the unit observing time beyond 10 to 20 seconds would filter out some of the real variations. Another way of lengthening the effective measurement for the real-time control operation is to decrease the control loop gain, as was done for observation number 10, where $g = 1/16$. This gain change was effected by changing a Right Arithmetic Shift instruction in the MODCOMP real-time software system. This gain change, by a factor of four, could have effected a factor of two decrease in phase variations if they behaved as independent variables, but only a 20% decrease was observed.

Observations numbers 5 and 6 were "drift" tests, where the VLA was calibrated on the selected radio source at the beginning of the time interval, after which phase control was turned off – the loop gain set to zero – and the phase variations observed as they drifted free over a 15 to 20 minute interval. Starting and ending rms phases listed in Table 1 indicate considerable growth.

Observation number 9 was inserted to test the scaling of source-strength with bandwidth, which was used to set the inferred source strength believed to be Voyager-comparable. The 2.6 Jy source via 3-MHz passband should result in approximately the same phase behavior as 0.5 Jy in 50 MHz. This it did, but the observation was inconclusive because the non-SNR-related effects appear to dominate the phase variations.

IV. Conclusion

Taken together, the several observations in the experiment demonstrate that one should be able to self-calibrate the VLA with a signal comparable in strength to that expected of Voyager at Neptune, without suffering serious SNR loss to the combined signal from the jitter in the real-time phase tracking process. Application of that result to the Voyager signal involved analogies in both the received signal frequency and the VLA processing bandwidths, so further testing of self-calibration with the Voyager spacecraft signal itself will be needed as the X-band capability begins to be assembled on the VLA. The phase variations induced at low elevations are a concern which needs further exploration, as Voyager will be at quite a low elevation throughout its encounter pass.

References

1. Napier, P. J., Thompson, A. R., and Ekers, R. D., "The Very Large Array: Design and Performance of a Modern Synthesis Radio Telescope," *IEEE Proceedings*, pp. 1295-1320, Nov. 1983.

2. Hjellming, R. M., Ed., *An Introduction to the NRAO Very Large Array*, 1983 edition, document available from NRAO to users and potential users of the VLA, (National Radio Astronomy Observatory, Socorro, NM 87801).

3. Layland, J. W. et al. "Interagency Array Study Report," *TDA Progress Report 42-74*, pp. 117-148, JPL, August 15, 1983.

Table 1. Summary of observations

OBS No.	Source	Flux	IAT	EL Angle	RMS Phase @ 2 min	Sig-Ph.	RMS Phase @ 30 sec	Notes
1	1803+784	2.52	4:12 – 4:21	45.1 – 45.3	2.5	3.7	–	Calibration
2	1714+219	0.53	4:25 – 4:37	77.8 – 77.5	3.2 ± 0.9	4.8 ± 1.4	6.5 ± 1.0	Ref @ Hi Elevation
3	2037–253	0.57	4:42 – 4:51	15.4 – 17.5	6.0 ± 2.0	9.0 ± 3.0	10.0 ± 2.6	Ref @ Lo Elevation
4	1831–126	0.14	4:57 – 5:09	41.8 – 42.6	3.8 ± 0.9	5.7 ± 1.4	6.3 ± 1.0	Nom. Vgr. Flux
5	1748–253	0.50	5:10 – 5:29	30.5 – 30.1	2 → 5	3 → 7	–	Drift Test No. 1
6	1748–253	0.50	5:30 – 5:47	30.1 – 29.5	4 → 10	6 → 15	–	Drift Test No. 2
7	1831–126	0.14	5:50 – 6:04	43.3 – 43.0	3.4 ± 0.3	4.2 ± 0.36	–	20 sec T_I
8	1831–126	0.14	6:05 – 6:21	43.0 – 42.4	4.2 ± 1.2	4.7 ± 1.3	–	60 sec T_I
9	1908–202	2.6	6:22 – 6:38	35.8 – 35.6	2.7 ± 0.3	4.0 ± 0.45	–	3 MHz Scale
10	1831–126	0.14	6:39 – 6:55	41.4 – 40.2	3.9 ± 0.9	4.1 ± 0.9	–	10 sec @ g = 1/16

Appendix

Extraction of Phase Jitter

Let the variance of the elementary phase estimates be denoted as V_{10}, and assume that this is due entirely to thermal noise or other independent noise sources. Assume further that V_{10} is small enough that nonlinearities can be ignored. These estimates are fed in real-time into the first-order phase-locking process with gain $g = 1/4$. The result of this process is that the individual local oscillators of the VLA antennas contain a refined estimate of the differential phase between the reference antenna and each of the others. Let $n_{i,t}$ denote the noise at time t on the phase-difference measured between antenna number i and the reference antenna, and let $\phi_{i,t}$ be the accumulated phase (error) in the antenna number i local oscillator relative to the reference antenna phase. Then assuming the nominal phase to be zero gives

$$\phi_{i,t} = \phi_{i,t-1} + g\,(n_{i,t-1} - \phi_{i,t-1})$$

$$= g \cdot \sum_{s=-\infty}^{t-1} (1-g)^{t-1-s} \cdot n_{i,s}$$

for the phase of the i^{th} local oscillator. The variance of the filtered local oscillator phases denoted V_{10f}, is

$$V_{10f} = V_{10} * \sum_{i=0}^{\infty} 9^2\,(1-g)^{2i}$$

$$V_{10f} = V_{10} * g/(2-g)$$

$$V_{10f} = V_{10}/7 \text{ for } g = 1/4$$

If the global solution is performed in final processing with the 10-second integrated cross products, the variance of those resultant phase measurements, denoted V_{m10}, is a combination of the residual phase offsets held by the local oscillators, and the additive noise during the measurement interval. In this case the additive noise of the current measurement interval is independent of the actual phase errors which depend only upon past measurement intervals. Let V_{10g} denote the noise variance on the global solution phases based upon a 10-second interval. Then by this independence,

$$V_{m10} = V_{10f} + V_{10g}$$

Although the noises on the 351 cross products are not independent, the individual noise components in the 27 correlated signal streams are. Hence, it is easy to show that the cross products are pair-wise uncorrelated, and therefore the covariance matrix of the inputs to the global least-squares process is diagonal with all non-zero entries equal to V_{10}. The formal covariance matrix of the 26 globally solved phases has all diagonal entries equal to $2{=}V_{10}/27$ and all non-diagonal entries equal to $V_{10}/27$.

Thus:

$$V_{10g} = \frac{2}{27} * V_{10}$$

Then, solving for the actual real-time phase errors, we find

$$V_{10f} = 0.79\, V_{m10}$$

Actual measurement intervals used are 30 sec and 120 sec, in order to reduce the amount of data to be handled to a more manageable level and to reduce the additive noise. As noise samples are independent between each 10-sec sub-interval, adding 'ℓ' of these together ($\ell = 3$ or 12) will reduce the noise variance by a factor of $1/\ell$ on all cross products (phase-difference measurements) *except* those involving the reference antenna. When phase-difference measurements involving the reference antenna are extended, the filtered local oscillator phase is anticorrelated with respect to the noise in the previous 10-second intervals.

Let $X_{i,t}$ denote the ℓ-interval average of the phase-difference measurement with respect to the reference antenna

$$X_{i,t} = \frac{1}{\ell} \sum_{j=0}^{\ell-1} D\,(i-r,\, t-j)$$

$$= \frac{1}{\ell} \sum_{j=0}^{\ell-1} \left[n_{i,t-j} - g \sum_{s=-\infty}^{t-j-1} (1-g)^{t-j-1-s} \cdot n_{i,s} \right]$$

Where $D(i - r, \cdot)$ denotes the 10-second phase-difference measurement, and $n_{i,s}$ as before denotes the additive noise on that measurement. By collecting like terms, and simplifying, the expression for $X_{i,t}$ becomes

$$X_{i,t} = \frac{1}{\ell} \sum_{j=0}^{\ell-1} n_{i,t-j} (1-g)^j$$

$$-g \cdot \frac{1}{\ell} \sum_{j=0}^{\ell-1} (1-g)^j \sum_{s=-\infty}^{t-\ell} (1-g)^{t-s} \cdot n_{i,s}$$

The variance of the $X_{i,t}$ is simple to compute as long as the $n_{i,s}$ is independent and becomes

$$V_x = V_{10} \frac{2 - 2(1-g)^{\ell}}{(2g - g^2)\,\ell^2}$$

For $\ell = 1$, this expression degenerates to the obvious $V_x = V_{10} + V_{10f}$. For $\ell > 1$, the V_x can be considerably less than V_{10}/ℓ, which is the contribution of the noise alone.

The ℓ-sample average of the local oscillator phases can be extracted from the above expression for the $X_{i,t}$ by removing the additive noise terms from the first summation. Denote the variance of this ℓ-sample averaged phase as $V\{\phi_\ell\}$. This variance is found to be

$$V\{\phi_\ell\} = \frac{V_{10}}{\ell^2} \left[\ell - \frac{2(1-g)\,[1-(1-g)^{\ell}]}{(2-g)\cdot g} \right]$$

For $\ell = 1$, this expression degenerates (as it should) to the expression for V_{10f}. Inserting numbers,

$$V\{\phi_\ell\} = 0.060 \cdot V_{10} \qquad \text{for } \ell = 12,\ g = 1/4$$

$$= 0.113 \cdot V_{10} \qquad \text{for } \ell = 3,\ g = 1/4$$

Global least-squares estimates of these ϕ_ℓ are extracted, and their statistics computed to derive V_{10f}. Of the 351 baselines used in the global solution, 325 contain additive noise terms which are uncorrelated from each other and from the ϕ_ℓ. The 26 baselines which are used in the real-time phase tracking process contain noise which is uncorrelated with the noise on the other baselines, but anticorrelated with the associated ϕ_ℓ. Since this apparent suppression of the added noise affects less than 8% of the baselines, and those only partially, we will *assume* in the following that this correlation between the real-time filtered phase estimates can be neglected, and solve for V_{10f} as determined from the accumulated measurements at 30 and 120 seconds. In retrospect, it is clear that the need for this assumption could have been avoided by discarding the reference antenna and the affected baselines from the global solution process. Using the assumption, we find:

$$V_{10f} = 1.14\, V_{m\,30} \qquad \text{using 30 second intervals}$$

$$V_{10f} = 2.27\, V_{m\,120} \qquad \text{using 120 second intervals}$$

A few other cases are also of interest, and can be extracted from the foregoing calculations: A) loop gain = 1/16, 10 sec real-time and 120 sec global integrals.:

$$V\{\phi_\ell\} = 0.786 * V_{10f}, \qquad \text{for 120 seconds, } g = 1/16.$$

Inverting as before,

$$V_{10f'} = 1.13 * V_{m\,120} \qquad \text{for } \ell = 12,\ g = 1/16$$

B) loop gain = 1/4, 20 sec real-time and 120 sec, global integral

$$V\{\phi_\ell\} = 0.619 * V_{20f} \qquad \text{for 120 sec, } g = 1/4$$

and

$$V_{20f} = 1.51 * V_{m\,120} \qquad \text{for } \ell = 6,\ g = 1/4$$

C) loop gain = 1/2, 60 sec real-time and 120 sec global integral

$$V\{\phi_\ell\} = 0.75 * V_{60f} \qquad \text{for 120 sec, } g = 1/2$$

$$V_{60f} = 1.24 * V_{m\,120} \qquad \text{for } \ell = 2,\ g = 1/2$$

These conversion factors will be used to convert from measured phase variations to an estimate of phase variations embedded within the real-time phase-lock process.