Third VLBA Design Review

JONATHAN D. ROMNEY
National Radio Astronomy Observatory
Charlottesville, Virginia

1985 December 10

INTRODUCTION

The Third VLBA Design Review workshop, held in Green Bank on 1985 September 10–11, addressed two primary objectives. The first third or so of the workshop was basically informational in character, aimed at maintaining a high level of awareness, among the many separate development groups, of plans and progress throughout the VLBA Project. Review talks from all major project areas outlined the current status of design and/or construction, emphasizing recent decisions and currently undecided questions; in some cases these presentations raised technical issues for later discussion. Much of the content of the reviews has been documented elsewhere (*e.g.*, the VLBA Book), and some is no longer current as a result of the workshop. For both these reasons, the present summary will not consider further this part of the meeting.

A second, and more critical, objective of the design review was to decide a number of pressing technical and budgetary questions. Previous face-to-face meetings of this kind have been invaluable in achieving a well-supported consensus on technical issues, and we attempted to do so in this review whenever possible. In addition, since the Project is now sufficiently advanced that fairly accurate cost estimates are available for all areas, this workshop had to deal seriously with budgetary issues for the first time.

As the review began the estimated cost to complete the VLBA showed a negative contingency exceeding a million dollars, due both to major increases in cost estimates for the station and array operations buildings and for the correlator, and to the adoption of attractive options in a number of project subsystems. The budgetary discussions thus focussed on establishing priorities and deferring implementation of some features of the Array until it becomes clear whether or not they can be afforded. These discussions centered on the VLBA Options List (VLBA Memo 473), and in particular on the proposed selection of cost-saving options considered by the VLBA Scientific Advisory Group (VLBA Memo 474), but considered other possibilities too. Following the convention of the Options List, all costs and savings in this document are expressed for the entire 10-station Array.

In practice, of course, technical and budgetary issues are closely linked, and decisions were based on considerations of both kinds. Therefore the following summaries are organized by topic, and follow neither the original program nor the actual sequence of sessions. I have taken some care to express accurately both the nature and, in controversial cases, the degree of support for any consensus reported. And since recent experience suggests that a number of controversial issues will be raised again, I have tried to reproduce for future reference the principal arguments for and against all major decisions. For simplicity, however, I have generally avoided attributing arguments to specific proponents. While necessarily subjective, my assessments are based on a careful review of recordings of the discussion sessions as well as on my notes made at the time.

Configuration and Site Selection

Owens Valley Site. A long-standing suggestion that the "Owens Valley Site" be located at some distance from Caltech's OVRO facility was discussed. VLBA Memo 482 (distributed at the workshop) considered alternative sites within Owens Valley and beyond, and recommended a location west of the Sierras in the Sacramento Valley near Fresno as the best of these. The main rationale for an alternative site is the formation of a short (and relatively sensitive) baseline between the VLBA site and the OVRO 130′ telescope when it is available; secondarily, any other location would also escape the unsatisfactory Owens Valley skyline.

On the negative side, the moderately bad RFI environment discovered at OVRO was expected to be worse at the alternative sites. And the availability of support for VLBA operations through OVRO (subject in any event to some debate) would not exist. In view of the infrequent and restricted availability of the 130′ telescope, there was in the end broad agreement to locate the VLBA site at or near OVRO.

Green Bank as the "Northeast Site". Previous requests to reconsider the VLBA configuration to locate a site at Green Bank were reiterated. The most cogent arguments in favor of this move were to exploit the National Radio Quiet Zone, especially in view of the unsatisfactory RFI environment found at Haystack and expected to extend throughout eastern Massachusetts, and to save site development and operation costs. The latter were summarized in VLBA Memo 486 (distributed at the meeting).

While there was general reluctance to sacrifice the aperture coverage by making major configuration changes, in fact expert opinion considered Green Bank "almost as good" a location as Haystack — but an unsatisfactory substitute for any other site. Despite protracted discussion, no agreement emerged on the extent to which significant savings in site development and operations could be expected. In partial resolution of this issue, Green Bank is to be considered an open possibility, as an alternative "northeast" site only.

Station Buildings

The architect/engineer's cost estimate for the station building represents a major escalation over earlier plans. It appeared imperative to achieve some savings in this area of the Project, and a variety of possibilities were considered.

RFI Shielding. Shielding to protect the front ends from RF radiation by digital circuitry in the station building is very expensive under present specifications, which call for certain performance standards. Since it is not clear at present whether shielding is necessary at all, one serious alternative would be to omit it from the first station building as an experiment; however, no reliable indication would be available from this first building before committing to the next two or three. Other possibilities discussed included: restricting shielding to the control room; the use of free-standing Faraday cages; and specifying suitable construction techniques, without performance requirements. While concrete decisions will have to await further negotiation with the architect and bids for construction of the Pie Town station building, a broad but reluctant consensus was that RFI shielding cannot be justified unless 500–750 k$ can be saved by some combination of these approaches.

Backup Power. Deleting the backup generators would leave the antennas unable to stow in the event commercial power fails. There was a universal disinclination to accept this risk. A lower level of saving was considered, retaining the generators but deleting the UPS's; since brief power interruptions are expected to be far more common than major failures, the small saving achieved by this was judged not to warrant the accompanying operational problems.

Building Size. Although no opinions were changed in the debate on this point, the majority supported the view that the station building as currently planned was "pretty generous" for an unattended facility. A smaller building should be considered, despite considerable skepticism about the architect/engineer's estimate of 1.5 k$ saving per square foot of building design.

RECEIVERS AND FEEDS

The VLBA's frequency coverage is quite dense, particularly in the range 1.7–22. GHz, and deletion or deferral of the receivers and feeds for one frequency was seriously considered as a cost-saving measure. The discussion turned on a number of interrelated aspects which are summarized separately below. In the end only one frequency, 2.3 GHz, appeared to be an appropriate negative option.

The general conclusions were: the 2.3 GHz receivers and feeds, and the associated 2.3/8.4 GHz dichroic reflector systems, should be *deferred* until later in the project; space should be reserved for this frequency and dual-frequency pair in the feed ring; and initial operation using the 5/15 GHz pair will be planned. Since this will compromise the sensitivity at 5 GHz, the concentration on ultimate centimeter-wave sensitivity should be shifted to the 8.4 GHz system. These measures should achieve savings of 600–800 k$.

A fractional version of this option — providing 2.3 GHz equipment only at perhaps 3 or 5 stations — is also a viable possibility, and input should be solicited from the geodetic community as to which VLBA sites were most critical as adjuncts to the existing geodynamic network.

The major points of discussion included the following:

Schedule. Development of many receivers and feeds is already well underway. The remainder fall into three groups: low frequencies, where the receivers and feeds are relatively inexpensive; high frequencies, which are scientifically crucial; and the 2.3/8.4 GHz pair. Further, the design of the feed ring must be completed in the very near future, which will commit us to the frequencies and pairings implemented therein.

Dual-frequency Capability. The requirement that the VLBA support *some* dual-frequency capability was emphasized repeatedly, and a variety of desirable pairings was enumerated: 2.3/8.4, 5/15, 5/22, 15/43, and 15/86 GHz. A distinction was drawn between two basic applications of this capability. One involves extending the coherence time of high-frequency observations by detecting fringes and monitoring phase variations at a lower frequency; this could probably be accomplished as well by reasonably rapid frequency switching. But the geodetic/astrometric application requires (nearly) simultaneous observations to calibrate the rapidly varying ionospheric phase contribution, which in turn requires use of dichroic reflector systems. It was noted that the 2.3/8.4 GHz dichroic systems (and indeed the 2.3 GHz feeds) would be the most expensive because of their size.

Uniform Frequency Coverage. Sufficiently dense coverage of the frequency axis is important because a number of astrophysical effects — synchrotron opacity and Faraday rotation foremost among them — depend sensitively on frequency. It was observed that omitting 2.3 GHz leaves an unusually large gap from 1.8 to 5 GHz; on the other hand, 8.4 and 10.7 GHz are unnecessarily close, and 15 to 22 GHz is also a relatively short step.

Compatibility. Of primary importance in this area is compatibility with the VLA, where 8.4 GHz will be available by the time VLBA observations begin. The 64-m DSN antennas operate at 2.3 in addition to 8.4 GHz; an inconclusive debate occurred as to whether these would be available for more than a negligible fraction of VLBA time.

LOCAL OSCILLATOR SYSTEM

Lock Points. The original specifications for the local oscillator system (VLBA Memo 303) called for tuning steps at $500n \pm \{100, 300\}$ MHz. The current design for the 2–16 GHz synthesizer, however, provides lock points only at $500n \pm 100$ MHz, and some concern was expressed that these steps were too coarse. The most problematic case appears to be at L band, where the receiver tuning range spans the greatest fractional bandwidth, but two different LO frequencies are required to access the extremes of this range; 43 GHz measurements, which will use L band as a first IF, may be affected as well. More generally, bandwidth synthesis observations spanning a wide band while simultaneously observing a lower frequency to remove ionospheric effects might be compromised.

While there was little dispute that at least a minor problem exists in this area, the preferred solution was not clear. It was agreed that the Electronics Group would investigate whether at the current stage of development the 2–16 GHz synthesizer design could be modified to incorporate more flexible tuning. Several objections were raised to changing the current rather simple design, however, and it was suggested that it might be simpler and less expensive in the end to provide additional fixed oscillators (as are already planned for S/X band operation) where necessary. In any case, the entire local-oscillator scheme will probably have to be re-evaluated in view of the changed plans for receivers, feeds, and dual-frequency pairing described above.

Fine Tuning. On a different frequency scale, some felt that the 10 kHz tuning steps available in the baseband converters were too coarse. To avoid the complexity and expense of a special "low fringe rate" mode in the correlator it might be preferable to dynamically offset the oscillators throughout the Array, but substantially finer frequency steps would be necessary to accommodate as many as 20 stations. Furthermore, the 10 kHz step is 16% of the 62.5 kHz bandwidth, which may produce unacceptably large baseline-dependent errors for narrowband spectroscopy. In the discussion of this point, it was agreed that the baseband converters should retain the planned tuning, primarily for reasons of cost efficiency (these modules are replicated many times) and to avoid delaying the construction schedule. If necessary, it may be possible to provide fine tuning by offsetting the reference frequency governing the first LO.

CHANNELIZATION, BASEBAND CONVERSION, AND RECORDING

Channel Bandwidth. The positive option to expand the maximum channel bandwidth to 16 MHz appeared to be realizable at very little cost, and was accepted with little discussion. It was noted, however, that exploiting this capability would require the correlator to handle 32 Msmp/s data streams. The proposal (VLBA Acquisition Memos 47, 48, and 51) that narrower bandwidths be implemented with active lowpass filters offered a saving of about 100 k$, and was also accepted.

Number of Channels. In view of the doubled maximum channel bandwidth, it was proposed that the number of baseband converters and channels be halved (to 8 and 16, respectively), thus maintaining the aggregate IF throughput intact. A lengthy discussion ensued, with those in favor of the proposal emphasizing the anticipated saving of 140 k$ (an estimate assuming the use of active filters), the reduction in archive volume and post-processing, and the disparity between IF throughput and current recording capacity. Opponents objected primarily to the reduced performance implied by the proposal in bandwidth synthesis, pulsar de-dispersion, low-frequency RFI avoidance, and Mark 3 compatibility. This debate was concluded more by exhaustion than by decision, but the proposal was generally judged a reasonable *initial* configuration; most participants found it important that the equipment involved be designed with sufficient modularity and flexibility to allow an eventual expansion in channel capacity and bandwidth by doubling to 16 baseband converters and 32 (16 MHz) channels.

Slow Playback. Two methods of achieving a reduced data playback speed were considered. Using the standard playback heads with separate equalizers a factor of 2 slowdown could be realized at rather little cost; factors as large as 8 could be achieved at much greater cost by a combination of extreme measures including very high-speed recording and special slow-playback heads. Both these options were rejected unanimously: the slowdown to half speed was judged insufficient to be worth even a small expenditure; the larger slowdown factor, while attractive, was too expensive.

CORRELATOR

Work on the correlator had already been suspended in order to maintain the antenna construction schedule despite budgetary constraints; current plans foresee a pause of about two years in the development of this subsystem. It was also noted that correlator cost estimates have suffered major increases over the original VLBA budgets. For these reasons, discussion of correlator issues was considerably less definitive than in most other project areas. The primary conclusion was to target a saving of 2.5 M$, without specifying in detail how this is to be achieved at present. Some limited changes in the specifications (described below) may nevertheless achieve significant cost reductions, and a further saving is to be expected from advances in microelectronics in the intervening two years. If necessary, deferral of some of the data-playback systems may be necessary. Besides these general conclusions, a strong consensus favored the following technical decisions.

Channelization. To parallel the changes planned for the baseband converters, the correlator should support 32 Mbit/s data streams if at all possible, but need only accommodate 8 input channels initially. The higher clock rate was judged likely to be feasible in the 2-micron CMOS semi-custom VLSI technology in which the current correlator design is implemented, although larger gate arrays might be necessitated for sufficient pipelining at this rate.

The number of input channels parallels the approach in VLBA Correlator Memo 41, where only half the maximum number of recorded channels can be correlated simultaneously; adding more channels has always been the preferred expansion path for this correlator architecture. These are the only changes foreseen in the correlator specifications, and current estimates anticipate a net saving of $\sim$150 k$.

Digital Filter & Accelerated Dump. Both these features are intended to support wide field-of-view observations. Arguments citing the scientific importance and post-processing load of these observations (and mentioning the relatively low cost of supporting them) prevailed over suggestions that one or both features be deleted at this time. It was pointed out that the proposed digital filter supports the slow fringe rotation case as well as widefield observations.

High-speed Archive Dump. The proposed European VLBI Network correlator includes the interesting capability of writing the correlation coefficients on a high-speed magnetic tape (*i.e.*, using a Honeywell 96 drive) after only a brief accumulation, which allows full "archival" over the entire primary beam, and later re-processing through the post-correlation backend hardware at offset phase centers. This capability was rejected for the VLBA correlator as being too similar to straightforward reprocessing in the short term (and too voluminous in the long term) to justify the cost of development; it would also require a post-correlation phase-shifting capacity which is not currently planned.

POST-PROCESSING

The intention of the Post-Processing Group to retain the total delay observable, in preference to model delay parameters and a residual, was universally endorsed. Nevertheless, there was also strong sentiment for carrying values of the most important components of the model delay as well;

while less fundamental, these would facilitate refining the overall delay model. Picosecond accuracy in these quantities was widely, but vaguely, seen as necessary to achieve the astrometric/geodetic potential of the VLBA. This and other questions regarding high-precision positional measurements will be addressed at a special-purpose workshop planned for next spring.

CALIBRATION TECHNIQUES

Amplitude Calibration. VLBA Memos 471 and 472 proposed a variety of methods for calibrating the amplitude of the interferometer visibility signal. One approach relies on monitoring the gain in the ALC loop, which would be determined by a precision attenuator (either mechanically or electrically switched). This method offers the most accurate calibration under ideal conditions, but is not suitable in the presence of drifting gain and also necessitates frequent absolute calibration against astronomical standards. Attenuators of the requisite accuracy are probably available at a reasonable cost.

Other methods require injection of calibration signals, either as incoherent noise pulses — possibly with a variable duty cycle — in an NAR mode, or as a continuous and stable coherent signal. The former is the conventional VLBI calibration; its drawbacks arise from potential interactions among the pulse signal, the phase-switching cycles, and the true interferometer fringes. The coherent calibration approach should be the most immune to adverse influences such as strong, impulsive interference, but a signal generator of satisfactory amplitude stability has not been demonstrated, and a substantial and somewhat risky development effort would be necessary.

After discussing all these methods, a strong feeling was evident that good practice, and our intention to construct the premier VLBI instrument, require a wide variety of calibration techniques, and that all three of these approaches should be seriously pursued. Sufficiently accurate attenuators will be incorporated in the baseband converters, and all receivers will be equipped for injection of calibration signals. There was some minority skepticism that development of a stable coherent signal generator was justifiable at this time. Future discussions of amplitude calibration should realize that the VLBA probably cannot effectively provide its own absolute flux scale, but will require occasional observations with other instruments, (*e.g.*, the VLA) to link its measurements to standard flux scales.

Phase Calibration. VLBA Memo 487 (distributed in draft form at the workshop) asserted that delay variations in the IF transmission cable could adequately be determined from the high-accuracy round-trip measurements planned for the parallel LO cable. Injection of a weak coherent RF signal for detection at baseband thus would be unnecessary; relative phase offsets among the baseband converters could be obtained from infrequent astronomical calibration observations. There was general agreement with the analysis leading to these conclusions. Nevertheless, since the Mark 3 experience has shown the coherent "phase cal" system to be a valuable diagnostic tool, and since in the immediately preceding discussion coherent signals had been approved for amplitude calibration, it was agreed without audible dissent again to plan for both these methods for phase calibration.

The diagnostic application of the coherent RF signals is enormously enhanced by having baseband detectors available at the stations, and indeed current plans for the VLBA station decoder module include such a function. It was not clear whether this would suffice, or if detectors would be required in the correlator, but since final design of the correlator has been delayed, this decision was deferred. It was pointed out that at non-VLBA stations, the stability of cable lengths and phase stability elsewhere along the signal path is likely not to meet VLBA specifications. Some provision will thus be required, either in VLBA-compatible local equipment or in the correlator, to measure delay variations for these stations.

Special Calibration for Solar Observations. A high-level noise signal for calibration of solar observations appears in the Options List as a positive option, and has in fact already been designed into some receivers and included in construction of the first units. VLBA Electronics Memo 30, however, pointed out that any fixed noise level will be unsuitable in view of the wide range and rapid variability of solar radio intensity, and recommended instead that the first of the methods discussed above under "amplitude calibration" be applied. Experience at the VLA has also shown solar calibration using high-level noise to be ineffective. It was decided, therefore, that no further "solar cal" noise sources would be built. The existing units will suffice to compare different calibration schemes during the Pie Town test period and verify that the alternative scheme is satisfactory.

MONITOR-AND-CONTROL SYSTEM

Monitor-and-Control Bus and Standard Interface. A number of prototype standard interface boards have been distributed to interested designers working in several VLBA subsystems. With the exception of the addressing scheme considered in the following paragraph, few comments were raised about the standard interface or the bus itself. There appeared to be general satisfaction with the interface's degree of accommodation to both "smart" and "dumb" devices. Current and potential users were urged to read the bus specifications (copies are available from the VLBA project office) and direct any questions or objections to the attention of the Monitor-and-Control Group.

Standard Interface Addressing Scheme. The standard MCB interface as currently specified assigns a unique address to each module, or more specifically to each interface board. The wish was expressed from many quarters for a system where the address could somehow be determined by the module's physical location, *i.e.*, by the plug through which it accesses the MCB, and possibly even for allowing the module to identify itself to the station control computer. Such an approach would enormously facilitate interchange of identical modules, and more generally was seen as an important contribution to simplicity and reliability of maintenance. After lengthy discussion of several alternatives, it was generally agreed that these advantages were worth the acknowledged additional design effort and increased cost of the interface, and a solid majority inclined toward a scheme using some pins of the standard interface board's connector to specify part of the address. This approach will be considered and refined further by the Monitor-and-Control Group.

Local Control of Equipment. No consensus at all could be reached on the issue of overriding the normal central control of equipment at the VLBA stations. The debate centered on whether such control should be exercised through the station computer, using a portable terminal or other I/O device, or by switches and connectors on the modules themselves. The former approach allows simpler modules and more centralized control of module status at the cost of software and cabling complexity; the latter affords greater modularity, simplicity of maintenance, and independence of the station computer. Reestablishment of central control was another inconclusive issue.

Signal Distribution. (This item may not quite fit, but seems more appropriate here than elsewhere.) There was some discussion of the distribution of timing and reference signals, and a request from the Acquisition/Recorder Group that this be defined more thoroughly. In addition to 5 MHz, 1 pps, and a station timing clock of unspecified frequency, several more specialized signals must be considered, including 10 kHz for the baseband converter oscillators and the 32 MHz sample clock. Coordination of synchronous detection for the proposed variable duty-cycle NAR must also be arranged. The station block diagram will be updated to show these relationships clearly.

OPERATIONS

Size and Scope of the Operations Center. Currently, 2.5 M$ is budgeted for the Array Operations Center, and this appears to suffice for the VLBA's portion (28,000 square feet) of a joint VLBA/VLA operations building. The unit cost of $90 per square foot is consistent with the experience of NMIMT in a similar building presently under construction. Some possible cost-reduction measures were outlined in VLBA Memo 498 (presented in draft form at the workshop), including halving of the auditorium, conference room, and library areas, eliminating some services such as a photo lab, and reducing the size or number of offices. In all, it was estimated that a 16% cost reduction could be achieved in the VLBA portion of the joint building.

Since the overall area per employee in the current plan is comparable with the situation at other NRAO sites, there was little support for reducing the size of individual work spaces. The appropriate number of scientist's offices, however, as well as the need for most of the other facilities, was a major controversy. The essence of this debate was whether a research center or a purely technical operations facility should be planned; no conclusion emerged despite protracted discussion. Two partial compromises were suggested, however. First, since the building will probably have to be financed in stages anyway, a natural division could be followed by building separately the technical support facility, and a building for the scientists and visitors supporting development and advancement of the Array. And, in a partial departure from the "all under one roof" philosophy, it might be effective to assign maintenance of front-end electronics and cryogenics to Green Bank, where much of this equipment is being developed.

At the conclusion of this discussion, it was noted that the Operations Center budget had only recently been raised from 2.0 M$. The Operations Group was urged to prepare a realistic plan for a reduced facility at this original cost.

Early VLBA Operations. The several phases of early VLBA operations and the transition to joint VLA/VLBA operations were considered. It became clear, however, that at present only the next few years could usefully be discussed, and for these the following milestones were generally agreed to. By the third quarter of 1986, the array control computer should have arrived and will occupy rented space; system tests at Pie Town (described further below) will begin using the local station control computer; and operator training will occur in parallel with these tests. Actual remote operation of the Pie Town and (probably) Kitt Peak sites will be achieved by the third quarter of 1987. It should not be necessary to perturb the current (or any future) construction schedule for the correlator to accommodate early operations, as the Haystack and Caltech facilities will provide sufficient and essentially compatible processing capacity to support initial observations using VLBA stations.

Test Plan for Pie Town. Two major methods of measuring and verifying the surface accuracy of the first VLBA antenna were described in VLBA Memo 489 (distributed in draft form at the meeting). The most definitive test would be a direct RF aperture efficiency measurement, at 86 GHz (or higher). To perform these tests expeditiously, it would be necessary to borrow a receiver from Tucson, and this may indeed be possible. Holographic observations, either in single-dish or VLBI mode, would probably follow; it was not clear whether a satellite in geosynchronous orbit or astronomical water masers offered the preferable source of radiation.

COMPUTER ISSUES

Database Management Systems. A general-purpose commercial database management system for the VLBA has been under discussion for some time. VLBA Memo 469 presented a tutorial introduction and considered some of the unique technical aspects of our application which may influence the design of such a system; VLBA Memo 485 (distributed in draft form at the workshop) contained further tutorial material and suggested selection criteria for an eventual decision.

The fundamental challenge faced in designing and specifying a VLBA database management system is the relative incompatibility between two major application areas: the remote fault-detection and diagnosis requirement (outlined in VLBA Memos 278 and 396) which emphasizes infrequent, rapid, but fairly straightforward retrievals from a voluminous array of monitor data from the VLBA stations; and the wide variety of general applications — scheduling, program tracking, performance and weather monitoring, correlator control, calibration, tape inventory, *etc.* — which require facilities for extremely flexible access to a diverse, but "reasonably" sized body of information. Exacerbating this division was the need to decide almost immediately what to do with the monitor data in order to meet the scheduled procurement of the array control computer, while most of the general applications remain as yet only vaguely defined.

This discussion led to two separate and somewhat indefinite conclusions. First, the Monitor-and-Control Group will evaluate commercial database systems for their application only. Only a few of the available systems are likely to satisfy the requirements for supporting the monitor data, and the evaluation should require only several man-months. It was also recognized that satisfactory graphics facilities would likely require a separate software package. The Monitor-and-Control Group will also update and improve the estimate of monitor-data volume.

Secondly, the current investigations of database management systems for the general VLBA applications should continue. The basic selection criteria restricting further study to true relational systems compatible with the VMS operating system were, evidently, widely accepted. Further, there were also frequent suggestions to "get something" as a test bed for gaining experience and, possibly, developing a realistic model of our application for later tests of several systems under trial license; several skeptics maintained, however, that little would be learned of general relevance beyond the specific test system.

Post-Processing Computer Budget. The current cost estimate for post-processing computers — basically unchanged from the original VLBA "Red Book" — describes a hardware configuration of four VAX-11/780 processors. It was generally recognized that this was (and is) intended more as a general descriptive specification of the computing power which might be needed, rather than a literal intention. In this approach, advances in hardware capacity were expected roughly to balance the increased requirements arising from new or more heavily used algorithms. Nevertheless, some suggestions were raised that an up-to-date plan should be maintained, based on current estimates of required computing power and hardware cost factors. Countering arguments maintained that computing "requirements" were extremely elastic and impossible to estimate accurately, and objected further to the effort involved, since procurement of the post-processing computers is still some time in the future, and a number of revisions would be required. At present a major revision seemed particularly inappropriate in view of NRAO's several plans for new hardware, which may change the outlook for VLBA post-processing. No further decision was reached on this point.

DÉJÀ VU

We reconsidered two major questions which had been considered at several previous meetings, and which most of us considered already settled. In both cases, the original decision had favored a more conventional solution over a less expensive, innovative, but technically risky and/or less compatible approach. At the urgent request of some true believers (among them some of NRAO's best engineering talent), I reluctantly reopened these issues to explore the potential savings and the ramifications of some technical points not previously realized.

Fringe Tracking at the Antennas. Implementation of fringe tracking at the individual antennas rather than in the correlator was discussed extensively (VLBA Memos 321, 326–329, 333, 334, 338, 345, and 349), and was the subject of an ad-hoc meeting, in the summer of 1984. At that meeting we decided against this approach for a variety of reasons: there were strong misgivings that the necessary phase accuracy could be achieved to avoid compromising astrometric/geodetic measurements; multi-pass processing of offset phase centers would be precluded or limited; rejection of receiver image bands would be reduced; and compatibility with non-VLBA stations (using Mark 3 equipment) would be abandoned. While none of these reasons was individually thoroughly compelling, in combination we felt they outweighed the primary advantage of potential cost reduction in the correlator. Indeed, since one of the correlator design concepts under consideration at that time was based on an existing VLSI element already incorporating fringe tracking, the cost reduction might well have been negligible.

Several more recent developments provided the introduction to the renewed discussion of this subject. First, VLBA Memo 432 pointed out the inevitability of beats between baseline fringe frequencies and the phase-switching cycles, unless fringe tracking *precedes* phase switching; this problem is particularly acute for the wide range of baselines in the VLBA. It was suggested that phase switching based on higher-order Walsh functions, expressible as products of square waves, would alleviate this difficulty, although the maximum sequency would thereby be increased. The proponents of using these functions were asked to investigate this possibility further.

Secondly, the compatibility objections have evaporated to a great extent, for several reasons. Informal polls of European observatories have shown a general intention to procure VLBA-compatible signal-processing and recording equipment, and the overall stretchout in the VLBA construction schedule has allowed more time for these and other non-VLBA institutions to do so. The many improvements now embodied in the VLBA design clearly serve as a strong incentive for this trend.

Finally, since cost reductions in the correlator subsystem are now a major priority, the saving from eliminating the fringe rotators, delay tracking logic, and half the multiplier/accumulator hardware in the correlator (directly realizable under the current design) is a serious consideration. Some of this saving would be consumed by the necessity for frequent accumulator dumps and signal processing to accommodate offset phase centers — all astronomers present concurred on the scientific indispensability of this feature — but nevertheless a reduction of $\sim$1–2 M$ in correlator cost was estimated.

In addition to these relatively new considerations, other advantages cited in previous discussions of fringe tracking at the antennas were reviewed. Baseline-dependent bandpass mismatches would vanish, as would the sensitivity losses (and unpleasant corrections) associated with digital fringe rotation and fractional-bit errors.

We then considered how this alternative approach might be implemented. Technically, the most suitable point to insert the required modulation would have been in the baseband converter oscillators. However, the expense of this modification (estimated variously at 2–5 k$) in each of 8 (or 16) modules per station would be high, and variable-phase sampling would then require additional, separate hardware. Most critically, the construction schedule for the baseband converters cannot tolerate a delay for the necessary design revisions. Accordingly, this implementation was not considered further, and discussion centered on the more elegant but difficult modulation of the station frequency standard's reference signal. Two possibilities were mentioned: a phase shift introduced somewhere in the maser receiver's phase-lock system, and an external offset applied to the standard 100 MHz signal.

To achieve sufficient phase accuracy, the modulation would be subject to quite stringent specifications; preliminary suggestions were phase equivalent to 1 picosecond delay, computed for an epoch accurate to 1 microsecond. Phase referencing between sources would require the offset phase to slew smoothly between source positions, without jumps. These specifications would ensure *random* errors sufficiently small for high-precision positional work, but concern about several potential systematic errors (described in some of the memos mentioned earlier) must still be answered.

With growing enthusiasm for this approach, we concluded the discussion by agreeing that the potential gains justified a serious further investigation. A preliminary study would require 3 to 6 months; if successful, a design and development effort amounting to 2 work-years will follow. The focus of both stages of the investigation should be the overall cost to the Array, which in the absence of strong technical grounds must control the eventual decision.

3- *vs.* 2- & 4-Level Quantization. Quantization of the sampled baseband signals had also been the subject of extensive prior discussion, although generally not in the VLBA "literature" — VLBA Memos 332, 421, 427, and 430 give some idea of the debate — and was decided in favor of 2/4-level quantization at the previous Design Review in January 1985. In the absence of new input we nevertheless managed another lengthy, and perhaps final, discussion.

The 2/4-level scheme, it is expected, will be used primarily in 4-level mode. This offers the advantages of the best sensitivity among all modes considered for band-limited observations, and the narrowest bandwidth requirement for a given continuum sensitivity. (The latter implies minimum archive volume and post-processing load, and most effective interference avoidance.) The 2-level mode is available as a simple subset, and will probably be used for special cases requiring gain independence, for geodetic/astrometric observations, and for compatibility with the Mark 3 system; as described above, this compatibility has become less critical a requirement than it once was.

The 3-level scheme achieves the most efficient use of tape (assuming that the 5-samples-in-8-bits "efficient" encoding scheme is employed) for a given sensitivity, and requires a substantially simpler multiplier table in the correlator. And since only one mode is supported, a general simplification in the quantization, formatting, and correlation hardware is also realized.

Trading off the merits of the two approaches on a cost basis is all but impossible because the relative advantages are virtually unquantifiable. The sensitivity enhancement in the 4-level spectroscopic observations, 7% by one standard of comparison, was regarded as significant because it cannot be obtained otherwise (except by observing longer — or enlarging the antennas). The 3-level scheme's 7% improvement in tape efficiency, on the other hand, was considered much less valuable, because the primary tape inefficiency will be induced by the gross quantization of daily tape changes. An old estimate of the extra cost of the 4-level multiplier under the present correlator design, $\sim$200 k$, was recalled, although it was suggested that the differential might be greater in a correlator based on PAL's. And while it was universally agreed that 3-level quantization was unattractive without the 5/8 encoding scheme, implementing this in the formatter was variously claimed to be "(not) much more expensive"; at a minimum it would impose an undesirable delay on design and construction of this hardware. Finally, although the importance of compatibility with existing 2-level systems may have declined, there was a reluctance to abandon it entirely.

At the conclusion of this debate, a substantial consensus was evident favoring the enhanced sensitivity and flexibility of the 2/4-level scheme. There did not seem to be a major, clearly definable cost factor. We ratified, again, the decision reached at the last Design Review to adopt 2/4-level quantization.