The EVN/Mark IV Station Unit Requirements

MARK IV MEMO #140
Revision D

930301

Contents

1 Introduction

A Station Unit (SU) is an element of a very-long-baseline interferometry (VLBI) system. In a VLBI system, multiple channels of digitised data are encoded on to multiple tracks on one-inch-wide instrumentation tapes. The primary function of the SU is to decode the data from these tracks and restore them into channel-based sample streams as if they had never been recorded. To perform this role satisfactorily, the SU needs logically to be tightly integrated with the playback recorder.

This document outlines what is required of the SU. The two major data formats with which the SU must maintain compatibility are very flexible and the ways in which these systems might be used are not completely specified. The references [?, ?] are the primary source of information about these formats. It is recognised that this flexibility may lead to questions by any company tendering for the development contract. Accordingly, the Project Director has nominated a Project Leader for the SU who will be available to resolve quickly any queries that might arise.

In order to capitalise on previous experience and to help potential tenderers to estimate the complexity of the system, a *straw-man design* has been undertaken [?]. This is available to all tenderers. The straw-man design is by no means exhaustive nor is any tenderer required to use it as the basis of their design. It is intended to be helpful and illustrative. It is recognised and expected that potential tenderers will want to negotiate with the project director detailed aspects of the development contract.

2 Background Information

2.1 VLBI Systems

The EVN/Mark IV System is a fourth-generation system for performing very-long-baseline interferometry (VLBI) using tape recordings of signals from several widely-spaced radio telescopes. The astronomical signals from each telescope are digitised and recorded on tape. The tapes are subsequently brought to a common site (the *correlator*) where they are correlated.

There are very stringent requirements for the maintenance of the coherence of the signals: almost all the frequencies used in the systems at each telescope are derived from a very stable source, usually a hydrogen-maser standard, and the data samples are effectively tagged with arrival times using very accurate clocks. The telescopes are always in motion with respect to the radio source. To maintain coherence, the signals presented to the correlator need to be from a common reference time-frame. The playback/correlation system needs to compensate for variations in the arrival times of signals at telescopes and also to compensate for speed fluctuations in the tape drives.

The radio signals are separated into a number, currently between 1 and 32, of frequency bands (*channels*) which are sampled at the Nyquist rate and digitised to 1-bit (sign only) or 2-bit (sign and magnitude) precision. The frequency bands can be between 16 and 0.0625 MHz wide in 9 octave steps.

The EVN/Mark IV playback system is required to handle Mark IV, VLBA and Mark IIIA formats, and to be extendable to double the bit-rate capacity of the standard Mark IV system by using 4 headstacks for 32 channels. The formats differ mainly in the way frequency channels are mapped onto recorder tracks.

For the Mark IIIA, which is mainly a subset of Mark IV, there is a one-to-one correspondence between channels and tracks: data are 1-bit digitised at rates up to 8 Mbit s^{-1} and recorded on a tape track. The full Mark IIIA format (Mode A) allows up to 28 channels to be recorded on 28 tape tracks. The standard Mark IV format records and outputs up to 16 channels of data. The EVN/Mark IV SU may thus be restricted to handling, at maximum in a single tape pass, Mark IIIA data in modes B or C which record 14 channels on 14 tape tracks.

For the VLBA system, data can be 1 or 2-bit digitised and there is no direct channel-to-track correspondence. The sample-rate from the digitisers can be as high as 32 Msample s^{-1} but the track signal rate is limited to 8 Mbit s^{-1}. Data from 1 channel can therefore be spread across several tracks for wideband channels or several channels can be multiplexed onto 1 track in narrow channel modes. As an additional complication, track signals can be switched dynamically between tracks – a practice known as *barrel-rolling*. The number of tracks recorded at a particular time varies with demand and many passes can be made over a tape.

Data handling for the Mark IV is similar to that of the VLBA except that more channels can be handled, multiple channels cannot be mutiplexed on to one track, more recorder tracks can be used and the maximum tape speed is increased so that the maximum signal rate per track is 16 Mbits s^{-1}.

In all 3 systems, data are organised into *track frames*. A track frame is a number between 2000 and 4000 of 9-bit bytes. Each 9-bit byte consists of 8 bits of information and a parity bit. A track frame starts with a header containing some ancillary data, a bit pattern used for frame synchronisation, a reference time code for the frame, and a CRC over the header. Data are coded differently in header bytes for VLBA compared to the Mark IV and IIIA which are identical. The rest of the frame consists of 9-bit bytes of 8 bits of packed signal samples and a parity bit. These frames are sent sequentially and bit-serially to tracks at rates up to 18 Mbit s^{-1} for Mark IV and 9 Mbit s^{-1} for VLBA and Mark IIIA systems respectively.

For the Mark IIIA and IV formats, the headers *replace* data in a track frame; for the VLBA which implements *non-data replacement*, the data are squeezed aside to make room for the headers. The correlator can ignore headers in Mark IIIA and IV modes but must see only unsqueezed data in VLBA mode. In both Mark IV and VLBA systems, the 8-bit data in the data fields of track frames may be modulated by a pseudo-random sequence synchronised to the frames. In addition, the Mark IV and VLBA systems permit completely flexible assignments of tracks to track-frame generators.

The VLBA system can record up to 36 tracks simultaneously, 32 of which are data tracks and the other 4 can be used to copy or replace other tracks or to carry cross-track parity bits for blocks of 8 tracks. In the latter case, when error rates are low or confined mainly to a bad track, it is possible to use these extra tracks to recover bad data at playback. The Mark IV can record 64 tracks simultaneously in *standard* mode and 128 in *extended* mode.

A general feature of both the Mark IV and VLBA formats is the flexibility and the redundancy built in. For example, it is often possible to record the same basic data in several different ways: with many tracks recorded at slow bit rates through to few tracks recorded at high bit rates. The VLBA system is barely operational and the Mark IV is not yet implemented. In both, many parameters such as the number of bytes per frame are programmable and wide use is made or will be made of field-programmable logic such as the Xilinx range of devices. This will allow some changes of format to be made if thought desirable. It is important that the EVN/Mark IV station unit have a similar flexibility through programmable parameters and the use of field-programmable devices.

2.2 VLBI Operations

2.2.1 Recording

A tape may be recorded at a telescope in a number of *passes* in which only some of the tracks are recorded at any given time. A tape pass can contain continuous observations (*scans*) of one or more radio sources each of which can last from a minimum time of about 1 second up to the duration of a whole pass or even multiple passes/tapes. The tapes recorded at different telescopes are not necessarily identical. For example, tape changes may not be made simultaneously at all telescopes because astronomical sources have different periods of visibility at different observatories. Weather problems and operator mistakes can also alter even the best planned usage of tapes.

Sometimes the tapes may be kept moving between scans so that the times in track headers are continuous, and sometimes the tapes may be stopped between scans to conserve tape. In the former case, some non-correlatable data may be recorded since the telescopes take finite times to slew between different radio sources, typically several seconds or more. In the latter case, track header times will be discontinuous at the boundaries between scans.

The choices available for recording mode are so complex and varied that the record process is almost always controlled by computer. However, human intervention is required for some aspects of recording so deviations from the desired program are bound to occur. Many of these deviations will be unrecoverable but some, such as late tape changes, must not result in more than the minimum loss of data. An electronic *log* is kept at each telescope which records what has been done. This log will be available at the processor so that it can be used to set up and control the station unit.

2.2.2 Playback

Tapes will not necessarily be played back at the speeds at which they were recorded. In general, tapes will be played back at the fastest rate for which data from the tapes can be processed. The limits are 18 Mbit s^{-1} for track bit rate and 32 Msample s^{-1} for correlator sample rate, whichever is most restrictive. Tapes may be played back more slowly than these limits dictate if it is required for the effective lag range of the correlator to be extended.

It is envisaged that tapes will be processed at the correlator in as efficient a manner as possible so that any data worth correlating will be correlated and, if possible, scans will be processed one after the other without stopping the tapes. The loss of a few seconds worth of correlatable data is undesirable but not, however, catastrophic.

It is necessary for station unit synchronisation and re-synchronisation to take place as quickly as possible. Continuously recorded tapes should not cause too many problems in crossing scan boundaries: track headers times will be continuous and the delay model (via the polynomial) cannot change by more than a few tens of milliseconds, at least for terrestrial baselines. The RAMs built into the station unit may even be able to absorb some of these changes without requiring sudden, gross changes in tape positions.

It is envisaged that the formats of individual tapes will be known from the telescope logs. The SU will typically read format information from track headers to verify the accuracy of the log rather than try to de-format the data from information in the track headers.

3 Requirements

3.1 General Requirements

The playback side of a EVN/Mark IV system consists essentially of a bank of recorders each equipped with a SU and a large, multi-telescope, multi-channel cross-correlator. A *Data Distributor* (a cross-point switch) may be placed between the SUs and the correlator if assignment flexibility is required. The whole system is controlled by the *Correlator Control Computer*, C^3, which sends commands via local-area network (LAN) to all other units. See Figure 1.

The SUs are required to convert data from tape tracks to channel sample streams synchronised to a reference clock keeping *wall-clock time*, to provide each sample with a *validity* bit which indicates that the sample is probably not in error, and to provide statistics about playback error rates. A more specific list of requirements is to:

- accept data from tapes recorded in Mark IV, Mark IIIA or VLBA formats from each of 72 tape tracks from 2 headstacks constituting up to 16 channels of data at a maximum bit rate of 18 Mbit s^{-1} $track^{-1}$.

- output to the correlator up to 16 channels of re-constituted channel samples (2-bit data and a validity bit) at a maximum rate of 32 Msample s^{-1}. In the widest band VLBA mode, it will be necessary to halve the tape speed (to VLBA speed) to satisfy this requirement. Other, slower, tape speeds may also be necessary to permit *correlator re-use* which effectively extends the lag range available in the correlator. The SU must be able to accept a range of track data rates in octave steps between 1.125 and 18 Mbit s^{-1}.

- permit total bit-rate capacity to be doubled compared to Mark IV by plugging extra cards into the SU, using 4 headstacks rather than 2 and by using 128 (out of 144) tracks at up to 18 Mbit s^{-1} $track^{-1}$ for double the number of output channels, i.e.32.

- de-assert validity bits during parity failures, during whole track frames which are improperly synchronised, and when header data pass to the correlator in Mark IIIA and IV modes.

- log and send to C^3 track-error-rate statistics.

- unsqueeze VLBA data recorded in non-data replacement mode.

- form wall-clock time by counting SYSCLK with respect to system second ticks, SYSSEC, and a *second* count distributed by LAN to each SU.

- output data samples synchronously with a system clock, SYSCLK, at up to 32 MHz. At a particular wall-clock time, the output data delivered should have been sampled at that time according to the telescope clock but offset by a delay computed by a Station Unit Control Computer (SUCC). The SUCC interpolates a polynomial which models the motion in time of the telescope. The polynomial coefficients are computed by the C^3 and distributed to the SUCCs via the LAN.

- keep the notional time-tagging of data accurate to half a SYSCLK period. It will be necessary to either repeat or skip particular samples to achieve this. This process must be controlled in hardware by a digital accumulator which can be loaded with an initial *value*, the delay-error (32-bits), and incremented by a *rate*, the delay-error rate (18 bits), for every SYSCLK transition.

- generate a signal, BOCF, (from second ticks, SYSCLK, and wall-clock time) which indicates the beginning of a *correlator frame*. Beginning at BOCF, the SU must overwrite the output validity bits by headers with about 224 bits containing several parameters for the *next* correlator frame such as: a frame serial number (32 bits), the delay-error (32 bits), the delay-error rate (18 bits), and individually for each channel, a phase (32 bits), a phase rate (32 bits), a phase acceleration (20 bits), a checksum for the header (16 bits). The delay and phase parameters must be computed by the SUCC from the delay polynomial, and from a phase polynomial for each channel. The period for which BOCF is asserted must be long enough to cover the period during which the correlator frame headers are transmitted.

- control the tape recorder by sending commands in VLBA protocol [**?**, **?**] through an RS232 connection.

- accept commands from, and generally communicate with, the Correlator Control Computer over an Ethernet local area network (LAN) using TCP/IP protocol.

- extract from each of the 16 output channels the quadrature components of at least two embedded sine waves. At a minimum, the allowable frequencies must include all multiples of 10 KHz up to 16 MHz.

It is desirable, but not required, that the SU be divided logically and physically into two separate parts: (1) a recorder-dependent part which would reconstruct synchronous channel bit-streams with time markers, and (2) a signal-delay-dependent part which would delay the signals by the appropriate amounts. The physical interface between the two parts would need to be accessible, perhaps on a backplane. Such a separation would allow the delay-dependent part to be incorporated in future systems using different recording technology and formats e.g. the S2 system.

More specific requirements are listed in the next subsection.

3.2 Specific Requirements

Inputs:

- **Data:** 72 differential pairs of ECL-level digitised tape signals in NRZM format, one for each of 72 tape tracks constituting up to 16 channels of data, at a maximum bit rate of 18 Mbit/s/trk. (including parity).
- **Clock:** SYSCLK at 32 MHz.
- **Tick:** SYSSEC at 1 pps.

Outputs:

- **Data:** 16 channels consisting of sign, magnitude and validity bits at a maximum rate of 32 Msample s^{-1} each. Note that data are digitised as 2's-complement but need to be converted to true sign and magnitude in the SU. For data digitised to only 1 bit, that bit will be the sign bit and the magnitude bit will be forced to the 1 state. The outputs will be in parallel as TTL levels. A vacant slot must be left in the SU crate for a card which will take these parallel outputs and transmit them bit-serially to the correlator by using high speed chips such as the Hewlett-Packard HDMP-1000.
- **Beginning of Correlator Frame:** BOCF – a pulse starting a correlator frame of duration between $\geq$ 10 millisec and $<$ 1 second. The signal should be active-high and be asserted for the 224-bit long header period in which delay and phase parameters overwrite the validity bits. A provisional list of what the headers contain is: initial delay-error (32 bits), delay-rate (18 bits), initial phase (32 bits), initial phase-rate (32 bits), phase acceleration (20 bits), frame serial number (32 bits), user bookkeeping (32 bits), correlator-frame invalid bit (1 bit), checksum (16 bits), etc.

Control:

- The station unit must be controlled by a VMEBUS controller with sufficient power to handle the double-length floating-point operations involved in polynomial interpolation, communications with the correlator control computer via Ethernet using TCP/IP protocol and the tape recorder via RS232 using VLBA Recorder Control Protocol [**?**, **?**], and control of the station unit.
- Some of the necessary commands from C^3 are
  - **Current Scan Interval:** define start and stop times for valid data in current scan.

- **Current Scan Delay Polynomial:** polynomial coefficients for delay as a function of *real* time. This might involve as much as a quadratic polynomial with 64-bit coefficients. (The VLBA uses a quintic spline polynomial calculated every 2 minutes to which linear approximations lasting 4 milliseconds are generated.) *Delay* in this context really refers to the desired times to be read back, to the sub data-sampling-clock level. It is a combination of geometric delay, clock offset, (perhaps) a difference between record and playback rates, and a large constant chunk resulting from the offset of the desired times from the current time.
- **Current Scan Phase Polynomials:** coefficients for a cubic phase polynomial appropriate for this station, to be used to generate the quadratic phase polynomials in every correlator frame header. There must be a separate polynomial for each individual channel.
- **Current Scan Epoch:** reference time defining the origin of delay polynomial.
- **Next Scan Interval:** as above, but for the *upcoming* scan.
- **Next Scan Delay Polynomial:** as above, but for the *upcoming* scan.
- **Next Scan Phase Polynomial:** as above, but for the *upcoming* scan.
- **Next Scan Epoch:** as above, but for the *upcoming* scan.
- **Segue Enable:** if set, make a seamless transition into the next scan by loading the next scan parameters when times beyond the current scan interval are encountered. If not set, stop the SU data stream (i.e. the tape drive) when times outside the interval are encountered. This command provides a mechanism for multiple (sequential) polynomials to be supplied to the SU during a scan, if necessary.
- **Set Time:** informs the SU what the real time will be at the next second tick.
- **Set Length of Correlator Frame:**
- **Prepare to Sync:** requests SU to find the current scan on tape, peak if necessary, verify formatter ID code, etc. The SU will report back when it is properly positioned with the actual time first available on the tape, or error codes in the event of failure.
- **Enable Sync Mode:** whenever true, the SU will attempt to drive the bit-error magnitude to less than 0.5 bit, and to stay in that state. Whenever the output data are more than 0.5 bit away from the the desired sync point, the frame it is contained in will be marked invalid.
- **Transport Command:** pass a Mark IV playback drive command along to the transport. If the command is a monitor command, pass the response back to C^3.

Construction:

- The SU modules will be constructed on 9U by 280 – 400 mm VME printed circuit or multiwire boards and will be housed in a VME card-cage(s). It is possible that the SU may be installed in the tape drive cabinet.

Test and Verification:

The SU is part of a complex system in which recording faults may have occurred weeks before they can be detected. It is of the utmost importance that proper test and performance-verification facilities are incorporated in the SU design at a fundamental level. It is expected that the successful tenderer will suggest to, and negotiate with, the project director the necessary test facilities to incorporate, and on suitable factory and site acceptance tests. Some means of testing are suggested in the straw-man design [**?**].

Maintainability

The firmware should be modular and well documented. If possible, code should be in C or C++ conforming to ANSI standards. Due consideration should be given to the practicability of future modifications to the firmware.

Documentation

The design of the SU should be fully documented to the extent that it should possible for SUs to be manufactured and maintained by a third party without the need for any redesign. All non-editable text and drawings should be in *Postscript* format, editable drawings should be in *Autocad Level 10* or *DXF* formats, PCB artwork should be in *Gerber* format, schematics in *Electronic Design Interchange Format* (EDIF) and text documents should be in *WordPerfect* format.

References

[1] VLBA Specification A56000N003 (901116) "VLBA Longitudinal Track Format" J.D. Romney.

[2] Mark IV Memo #137 (930217) "The Mark IV Tape Format – Revision 1.0" M.I.T. Haystack Observatory.

[3] Mark IV Memo #141 "The EVN/Mark IV Station Unit: Straw-Man Design" Revision D, B. Anderson.

[4] VLBA Acquisition Memo #71 (revised 920522) "VLBA Acquisition Recorder Controller Communications Protocol" R. Cappallo.

[5] VLBA Acquisition Memo #238 (revised 920522) "Recorder Controller Commands" R. Cappallo.