

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

HAYSTACK OBSERVATORY

WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764

Fax: 781-981-0590

24 November 2004

TO: Mark 5B development group
 FROM: Alan R. Whitney and Roger J. Cappallo
 SUBJECT: Mark 5B design specifications

1. StreamStor disk format

Within each scan, the data on the SS disk are divided into equal-length ‘disk frames’ (DFs). Each DF carries a *frame header* of four 32-bit words followed by 2500 32-bit words of data. There is always a DF boundary at each UT second tick.

Recording starts on the first DOT second tick after a VSI-S ‘receive=start’ command is issued, so each scan always begin with a DF header. Recording stops on the first DOT second tick after a VSI-S ‘receive=stop’ command is issued.

1.1 DF format

The format of the non-data-replacement DF header is shown in Table 1; each DF header contains the following information:

Word 0 – A fixed synchronization word (value not yet determined).

Word 1

Bits 31-16: User specified (e.g. station ID)

Bit 15: T – tvg data

Bits 14-0: DF # within second (starts at zero on second tick)

Words 2-3 – VLBA BCD Time code and 16-bit CRCC

	Bit 31		Bit 0
Word 0	Sync word (TBD)		
Word 1	User-specified (16 bits)	T	Frame# within second (starting at 0)
Word 2	VLBA BCD Time Code Word 1 (‘JJSSSS’)		
Word 3	VLBA BCD Time Code Word 2 (‘.SSSS’ plus 16-bit CRCC)		

Table 1: Disk Frame Header format

2. Test Vector Generator

The TVG will operate in two modes (software selectable):

1. Standard VSI mode: Uses a standard VSI tvg sequence which restarts on every DOT second tick.
2. Incrementing mode: Generates an incrementing 32 bit binary pattern which restarts at zero on every DOT 100-second tick. This mode will be particularly useful in testing the delay management of the SU functions of the DOM.

3. PDATA Management

PDATA arriving on the 80-pin VSI-H DIM connector is shunted to a serial port on the host PC, where each PDATA message during a scan is recorded with a UT time-tag and acted on, as necessary, in accordance with the VSI-H specification. At the end of each scan, the set of time-tagged PDATA messages collected during that scan are written to the SS disk (in ASCII format) as a separate file and tagged in the SS directory as the associated 'PDATA file'.

On playback, the 'PDATA file' is read *prior* to starting actual data playback and held in the host PC. During data playback, the QDATA messages associated with each playback second tick is issued through the host-PC serial port, which will be connected into the 80-pin VSI-H DOM connector.

4. Station Unit Emulation

The DOM can be configured to emulate a Mark 4 correlator Station Unit, as shown in Figure 1. Define some term and signals:

TOT ('tape observe time') – the time scale on the recorded data; the 'TOT' acronym is a holdover from magnetic tape days!

TOT1PPS – TOT second tick, generated each time a DF Frame #0 is read from the data disks.

ROT ('reconstituted observe time') – the correlation processing time scale, which is pegged to an abstract clock at the center of the earth. Due to the fact that playback may be speeded-up or slowed-down by up to a factor of 16, ROT second ticks may occur at a rate from 1/16 to 16 times the record rate, though maximum playback bitstream rate is always 32MHz.

DOM1PPS – correlator second tick, which has the same rate as the correlator wall-clock (DPS1PPS).

4.1 Sequence of operations

The sequence for SU emulation is as follows:

Preparation:

1. The correlator software selects a particular DOM1PPS tick on which playback is to be started, which is assigned (by correlator software) to correspond to the ROT start second tick for scan processing, which we define as ROT₀.
2. In advance of the actual playback start, the Mark 5B software commands a pre-load of the DOM Data Buffer, which always starts precisely with a TOT second tick, defined as

TOT₀ (see footnote¹). The minimum size of the Data Buffer is 256MB, corresponding to two TOT seconds of data at 1024 Mbps ($f_{\text{CLOCK}}=32\text{MHz}$) or up to 32 TOT seconds for the slowest record rate ($f_{\text{CLOCK}}=2\text{MHz}$)². The Data Buffer will always be filled with this pre-load data. The Data Buffer may, of course, be filled at the maximum rate provided by the StreamStor disks.

3. In advance of playback start, the Mark 5B software loads the delay/CF-header data for the first CF to the DOM.

Playback:

1. The Mark 5B software issues a 'start playback' command to the DOM sometime in the DOM1PPS period prior to the DOM1PPS start tick (chosen by the software to correspond to the desired ROT start time). At the instant of the start DOM1PPS, the DOM starts playback of data from the Data Buffer according to the delay parameters and enables TOT1PPS and CFINT interrupts. The data are divided into CF's as specified, with the data dynamically delayed according to the delay model and CF headers being created as required (see Section 4.3).
2. A TOT1PPS interrupt is generated every time the DF count *should* roll over to zero (if a fill pattern obliterates a header, a TOT1PPS tick is still generated). The *whole number* of seconds in the VLBA time code is posted to a register and a TOTCOUNT register is incremented. At each TOT1PPS interrupt (or any other time) the software may check that the TOTCOUNT and posted VLBA time code are consistent.
3. At each CFINT interrupt, including the one occurring at the start of playback, a CFINT interrupt is generated which causes the Mark 5B software to load the delay/CF header data array for the next CF into the DOM. In this manner, the DOM always has a fresh set of delay/CF-header data available for each CF. The process continues through the duration of the scan. See Section 4.6 for more details.
4. When the DOM read point passes the midpoint of the Data Buffer, the DOM reads the next data from the disks into the first half of the Data Buffer and then pauses. When the DOM read point reaches the end of the Data Buffer, the read point is moved back to the beginning and disk data are backfilled into the second half of the Data Buffer, etc. The average read rate from the disks during the backfill operation should be regulated to be perhaps no more than about 10% than the playback rate in order to maximize the dynamic delay range during playback.
5. A 'pause playback' command may be issued at any time during playback, in which case the output continues until the end of the current ROT second; at this point all further

¹ The actual choice of TOT₀ depends on the sign of the initial delay at ROT₀, which will always have a value between +/-½ ROT second. If the initial delay is positive (as specified by the sign of the initial delay parameter), the DOM will begin playback at TOT₀+abs(initial delay); if the initial delay is negative, the DOM will begin playback at TOT₀+1-abs(initial delay). Note: Delays greater than ½ ROT second can be handled by correlator software by adjusting the TOT₀ value by integer seconds.

² Data are loaded as 32-bit words into the DOM Data Buffer (as if there are always 32 active bit streams); dummy data may be loaded into the parts of the 32-bit word that correspond to unused bit streams. Each 32-bit word is accompanied by a validity bit. If the scan ends before the Data Buffer is filled, the remainder of the Data Buffer are flagged invalid.

output data will be marked invalid, though BOCF and RCLOCK signals will continue to be active.³

6. A 'resume playback' command will re-start playback on the next DPS1PPS tick, picking up exactly where the data output was stopped with the 'pause playback' command, but with an updated delay model.
7. If the recording or playback configuration changes between scans, the DOM must be reset and the internal data buffer re-loaded before playback can be restarted.

Figure 1 shows a typical SU timeline for a 'pause/resume playback' scenario at a scan boundary. Figure 2 shows a typical SU timeline for a 'pause/re-start playback' scenario at a scan boundary.

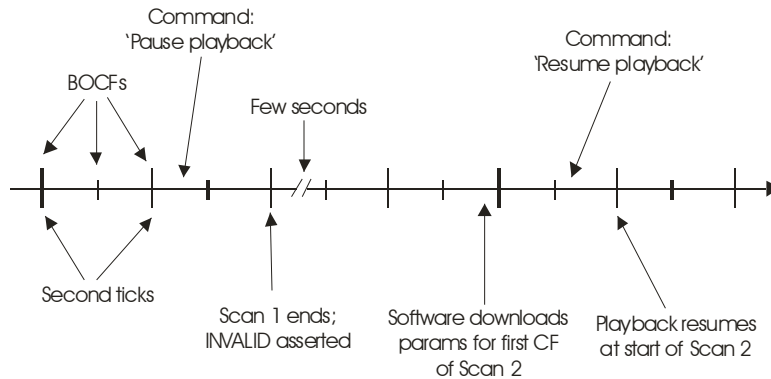


Figure 1: Typical SU timeline for 'pause/resume playback' at scan boundary

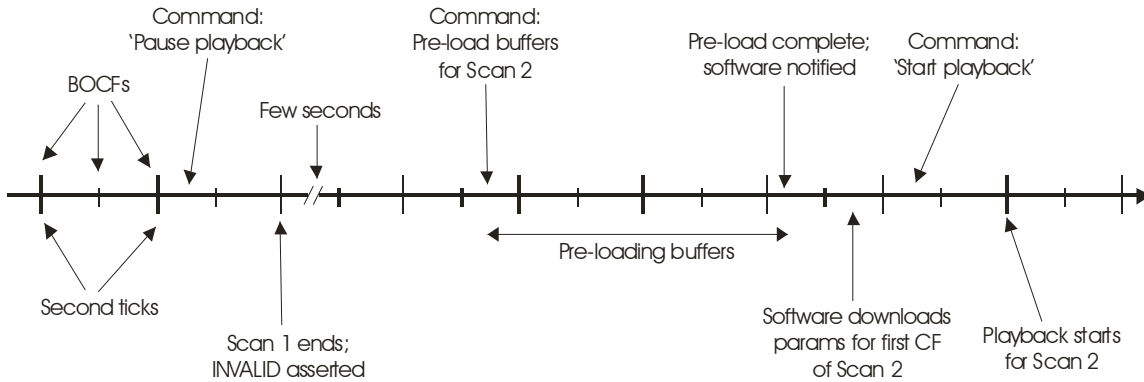


Figure 2: Typical SU timeline for 'pause/restart playback' at scan boundary

4.2 Sign, Magnitude Codes

When operating as an SU, the DOM must remap the Mark 4/VLBA according to Table 8, which is identical to the existing Mark 4 SU remapping.

³ If no 'pause playback' command is issued at the end of a scan, playback will continue unabated through subsequent scans. A ROT time discontinuity of exactly an integral number of seconds will take place at scan boundaries.

BBC analog output (mv)	Mark 4/VLBA code		DOM Output	
	Sign (S)	Mag (M)	Sign (S')	Mag (M')
>220	1	1	1	1
0→220	1	0	1	0
0→-220	0	1	0	0
<-220	0	0	0	1

Table 8: Codes used when Mark 5B is operating as an SU (2-bits/sample)

The code transformation in Table 8 may also be expressed as logically as $S' = S$, $M' = M \oplus \bar{S}$.

In the case where the channels are 1-bit/sample, the corresponding magnitude bit-stream must be dynamically constructed as $S' = S$ and $M' = I$.

4.3 Correlator Frames

When operating in SU mode, the output data are always organized into 16 2-bits/sample channels, where all sign bits appear on even-numbered bit streams and magnitude bits on the corresponding odd-numbered bit streams.

All output bit-streams are divided into Correlator Frames (CF), which have the following characteristics:

1. The length of a CF will be specified by the user in units of RCLOCK cycles, but the rate will not exceed 32 CFs per DOM1PPS period; the maximum length of a CF will be 0.5 DOM1PPS periods (corresponds to maximum of 1.6×10^7 RCLOCK cycles at 32 MHz). There are always integral number of CFs per DOM1PPS period.
2. A CF begins on every ROT second tick.
3. For each of the 16 'magnitude' bit streams, the 240 bits occurring during BOCF are replaced by CF header data provided by the host computer; the CF-header data are different for each 'magnitude' bit-stream.

4.4 RCLOCK

When operating in SU mode, RCLOCK always operates at f_{DPSCLOCK} (32 MHz), even though the DOM output bit-stream rate may be 2, 4, 8, 16 or 32 MHz. This is unlike normal DOM operation where f_{RCLOCK} is always the same as the DOM output sample rate. RCLOCK shall be continuous and uninterrupted while operating in Station Unit mode.

4.5 BOCF

A Beginning of Correlator Frame (BOCF) signal, which applies to all bit streams, is asserted at the beginning of each CF for a period of $240 * n$ RCLOCK cycles, where n is software selectable to be 1, 2, 4, 8 or 16. (Note that the RCLOCK rate is always constant, as stated above.) After setting of the BOCF period and length, the BOCF signal shall be continuous and uninterrupted from scan to scan until the BOCF is explicitly reset.

4.6 Delay Management

An initial delay of the data (TOT wrt ROT), which applies to all bit-streams, is specified for the beginning of each CF, as well as a constant delay-rate to be used by the SU during the CF. The maximum initial delay will be ± 0.5 seconds (TOT wrt ROT). During the course of a single CF, the delay may change by a maximum of 600 samples, determined by the delay parameters applied to the CF.

Three delay parameters are provided for every CF. Minimum requirements are:

1. Initial TOT-ROT offset: 24-bit magnitude + 1-bit sign
2. Initial fractional-bit delay: 32 bits
3. Delay rate: 18-bit magnitude + 1-bit sign

It may be simplest to manage these as three 32-bit words, in which case the maximum data rate (at 32 CF/second) for the delay parameters is $32 \times 12 = 384$ bytes/second.

During a CF, the delay will be managed by a delay-generator exactly analogous to the delay generator in the Correlator Chip (see paper on Mark 4 correlator by Whitney et al). At each carry out of the delay generator, a sample will either be dropped or duplicated, depending on the sign of the delay rate.

The delay parameters for the first CF must be provided to the DOM preceding the start of playback. Subsequently, a CFINT interrupt signal is generated at each CF boundary (including the one at the first instant of the scan), which will prompt the software to download the delay parameters for the *next* CF. It is the responsibility of the software to deliver the CF data in a timely manner.

Within a scan, the SU should be able to handle a delay step of up to about 20 msec with no loss of data synchronization and no invalid data.

4.7 Data-Checking and Valid-Data Management

During recording, the DIM replaces normal data being record to disk by a user-specified 'DIM fill pattern' whenever PVALID is enabled and de-asserted, or if software declares the data 'invalid'.

During playback, the DOM obeys the following rules in managing the QVALID signal:

1. De-assert QVALID for all samples associated with FPDP words which match a user-specified 'DIM fill pattern' or 'SS fill pattern'. During this period, keep track of where DF headers are supposed to be.
2. When fill data cease, re-assert QVALID, subject to rule 3.
3. If the correct DF header does not occur in the expected position, *all* subsequent data will be marked invalid. The DF checks are:
 - a. Playback must start precisely with a TOST DF frame header
 - b. All DF headers must occur in the correct position and with the expected DF sync and frame number (not necessary to check CRC).

If either of these checks fail, an error flag is posted and all subsequent data are marked invalid.

4.8 Normal VSI playback

During normal VSI playback from disk, the SU emulation features are turned off and the DOM acts in accordance with the VSI-H specification.

Notes:

1. RCLOCK and ROT1PPS will be continuous and uninterrupted so long as the playback configuration (f_{RCLOCK} , data multiplex ratio from disk, crossbar settings) does not change.
2. The 'pause playback' and 'resume playback' features will work identically as in SU mode.
3. The TOT1PPS is generated in exactly the same manner as in SU mode and the same information made available to the software (TOTCOUNT and whole number of seconds in the VLBA time code)
4. The ROT1PPS-to-R1PPS delay model may be updated once per ROT second in a manner entirely analogous to the SU. The VSI specification mandates only a fixed delay over a ROT second, but the same linear model used in SU mode may be used, if desired. The CFINT interrupt is replaced by an analogous ROTINT interrupt signal.

4.9 SU Output

The SU-emulation output will be on the standard 80-pin VSI output connector. For connection to a Mark 4 correlator, this output will go to a separate "correlator interface board" (CIB) which converts the data to high-speed-serial format compatible with the Mark 4 correlator input.

4.10 Phase-cal extraction and state counting

As shown in Figure 1, an output from the DOM feeds undelayed data, along with the necessary timing and validity signals, for state counting and extraction of the phase-cal signals. [Details of these functions are TBD.]

4.11 Functional limitations of this approach

There are two primary functional limitations of the above-outlined approach compared to the existing Station Units:

1. The Station Unit allows each channel to use an independent delay model within the limits of its on-board buffer memory. This capability is not supported on the Haystack, USNO or MPI correlators; support on the JIVE correlator is unknown. The approach outlined here allows only a single common delay model across all channels.
2. The Station Unit allows the Validity line associated with each channel to be controlled according to a pulsar-gating model. The approach outlined here has no pulsar-gating capability.

Adding either or both of these capabilities will significantly complicate the Mark 5B design. If either of these capabilities are needed, the Mark 5B data may be played back through a (to be) upgraded Mark 5A connected to a standard Mark 4 Station Unit.

4.12 Additional comments

With a direct e-VLBI connection into the correlator it may be difficult for the Mark 5B to both move data from the network to the SS card and to manage the SU functions at the same time due to congestion on the PCI bus. At the very least, it seems likely that the StreamStor card and the Mark 5B interface card must be on independent PCI buses. In any case, we need to keep this particular application in mind as the Mark 5B is designed.

5. PCI Interrupts

The PCI interrupt can be generated by the occurrence of up to five signals shown on Figure 1. An interrupt mask will allow the user to enable/disable each interrupt source individually. An interrupt register (may be separate registers for DIM and DOM, if desired) readable by the user specifies which interrupt source(s) caused an interrupt; reading this register will clear all bits in the register. Table 9 lists all the PCI interrupts, their purpose, and when they are enabled.

Name	Source	Description	When enabled/used?
DIM1PPS	DIM	External station 1PPS tick to which DOT Clock is synchronized	When DOT Clock is being set (normally happens once per power-up)
DOT1PPS	DIM	Internal DIM 1PPS tick	During most DIM operations; provides second tick marker to coordinate software operations.
DOM1PPS	DOM	External correlator 1PPS tick.	Prior to scan-playback start-up; provides second tick to coordinate playback startup; disabled during actual scan playback
TOT1PPS	DOM	Generated whenever DF Header #0 is read (corresponding to TOT second tick)	During scan playback. TOTCOUNT and VLBA integer-second count returned to user.
CFINT	DOM(SU)	Generated at every CF boundary	During scan playback
ROTINT	DOM(VSI)	Generated every ROT second tick	During scan playback

Table 9: PCI Interrupts

5.1 DIM1PPS Interrupt (DIM)

The DIM1PPS interrupt is generated at each tick of the selected DIM input second tick (1PPS, ALT1PPS, or ALT2PPS). It allows the Mark 5B software to initiate setting of the DOTCLOCK cleanly between DIM input second ticks.

5.2 DOT1PPS Interrupt (DIM)

The DOT1PPS interrupt is generated at each tick of the DOTCLOCK to coordinate DIM software operations.

5.3 DOM1PPS Interrupt (DOM)

The DOM1PPS interrupt is generated at each tick of the selected DOM input second tick (DPS1PPS or DPS1PPSX) to coordinate DOM software operations.

5.4 TOT1PPS Interrupt (DOM)

At playback start and continuing through the duration of a scan, a TOT1PPS interrupt is generated every time the DF count *should* be zero (if a fill pattern obliterates a header, a TOT1PPS tick is still generated). The *whole number* of seconds in the VLBA time code is posted to a register and a TOTCOUNT register is incremented. At each TOT1PPS interrupt (or any other time) the software may check that the TOTCOUNT and posted VLBA time code are consistent. If this check fails, the software will take appropriate action. Note that playback may, in some cases, be speeded up by a factor of up to 16, in which case the TOT1PPS interrupts will be generated at a rate of up to 16 per DOM1PPS period.

5.5 CFINT Interrupt (DOM SU mode)

The CF Interrupt Generator is initiated by playback start and generates a ‘CFINT’ interrupt at every CF boundary, including at the initiating DOM1PPS tick. The CFINT must be serviced by the host CPU before the following CF boundary by loading the delay/CF-header data for the next CF. The maximum rate of CFINT interrupts is 32 per DOT1PPS period. It is the responsibility of the software to deliver the CF data in a timely manner.

5.6 ROTINT Interrupt (DOM VSI mode)

The ROTINT is generated only during VSI playback and is analogous to CFINT in SU mode. The user response to ROTINT is to load the delay model for the next ROT second interval. If more convenient for the designer, the TOT1PPS interrupt could be used for this purpose in VSI playback mode.

6. Summary of Differences between VSI and SU playback modes

Table 10 summarizes the differences between VSI playback mode and SU playback mode.

Item	VSI playback mode	SU playback mode
f_{RCLOCK}	$f_{\text{RCLOCK}}=f_{\text{playback_rate}}$	$f_{\text{RCLOCK}}=f_{\text{DOMCLOCK}}$ (normally 32MHz) (i.e. RCLOCK rate is independent of output sample rate)
Number of active output bit streams	1-32; bit-stream meaning unspecified	Always 32, organized as 16 channels; sign bits on even-numbered bit-streams; mag bit on odd-numbered bit-streams; for sign-only channels, mag bit must be constructed.
Correlator Frame	Not defined or used	Data are segmented into 2 to 32 Correlator Frames per DOM1PPS.
CF Header	None	Replaces first 240 samples of the magnitude bit-stream for each of the 16 output channels
BOCF	Not defined or used	BOCF=1 for first $240 \cdot 2^n$ samples of each CF ($n=1,2,4,8$ or 16)
Delay/delay-rate	Fixed integer-sample delay can be specified for each ROT second	Independent SU delay/delay rate model specified for each CF
Phase-cal/state-count	None	Extracts up to 16 phase-cal tones from each of the 16 channels of data.

Table 10: Differences between VSI and SU playback modes

7. Compatibility with Playback on Mark 5A

The VLBA time-tags are included in the Mark 5B frame header at appropriate intervals so that VLBA-format 'track' data can be reconstructed by an upgraded Mark 5A playback system. Table 11 shows the VLBA-format playback modes that will be supported in such an upgraded Mark 5A system.

#Active DIM input bit-streams	Fanout ratio to VLBA-format tracks	#output VLBA tracks
32	1	32
32	2	64
16	1	16
16	2	32
8	1	8
8	2	16
4	1	4
4	2	8
2	2	4

Table 11: Playback 'VLBA-compatible' modes to be supported by upgraded Mark 5A

The choice of Mark 5B frame length of 2500 data-words/frame does not allow support of compatibility mode with a single active DIM bit stream, but this is not judged to be serious.

7. Correlator Interface Board

The Correlator Interface Board (CIB) transforms the DOM output, when operating in SU mode, into the form necessary for input into the Mark 4 correlator. Figure 4 shows a simplified block diagram of the CIB.

The CIF is 'dumb' in the sense that it needs no communication with the outside world.

Code Transformation

From the $\{S', M', QVALID\}$ DOM output for each channel, the CIB constructs an $\{S'', M'', V''\}$ output triple for each of the 16 channels according to the following rules:

- While $BOCF=1$, the CIB output is $\{S', M', M'\}$ (i.e. the station model parameters carried on the DOM magnitude streams are copied to V'' of each corresponding channel)
- For DOM output $\{S', M', 1\}$, the CIB output is $\{S', M', 1\}$.
- For DOM output $\{S', M', 0\}$, the CIB output is $\{0, 1, 0\}$; this prevents such samples from being confused as a 'valid zero' by the correlator chip.

Electrical Transformation

The electrical output of the CIB must be multiplexed and converted to high-speed serial interfaces compatible with the Mark 4 correlator.

8. Mark 5B Front-panel Status LEDs

8 tri-colored LED's will be located in the 2cm x 8cm cutout in the lower right corner of the front panel. The LED status panel will be implemented on a small PCB and connected to the I/O board via a small cable. Table 12 indicates the function of each LED and the meaning of each color when operating in DIM or DOM mode.

Function	DIM mode	DOM mode
Operating mode	G – DIM B – DOM/SU R – DOM/VSI Off – FPGA programming failure	
1PPS	Flash at 1PPS rate: R - <u>Not sync'ed</u> ; flash with selected external 1PPS G - <u>Sync'ed</u> ; flash with DOT1PPS	Flash G with DSP1PPS
Data validity	G – 'valid' data at DIM input B – TVG selected and active R – 'invalid' data at DIM input	G – 'valid' data at DOM output B – TVR active; TVG data recognized R – 'invalid' data at DOM output
FPDP data flow	G – writing valid data to FPDP R – writing fill pattern Off – no data flow	G – reading valid data from FPDP R – reading fill pattern Off – no data flow
Error	G – normal operation R - Internal error detected	
Delay buffer state	Off - NA	G – write in 2 nd quadrant behind read B – write in 3 rd quadrant behind read R – write in 1 st quadrant behind read Off – write in 4 th quadrant behind read
Software controlled	-	-
Software controlled	-	-

Table 12: Status LED functions (G-green, R-red, B-blue)

Suggested LED layout:

Mode	1PPS	Data Valid	FPDP
○	○	○	○
Error	Buffer	S/W1	S/W2
○	○	○	○

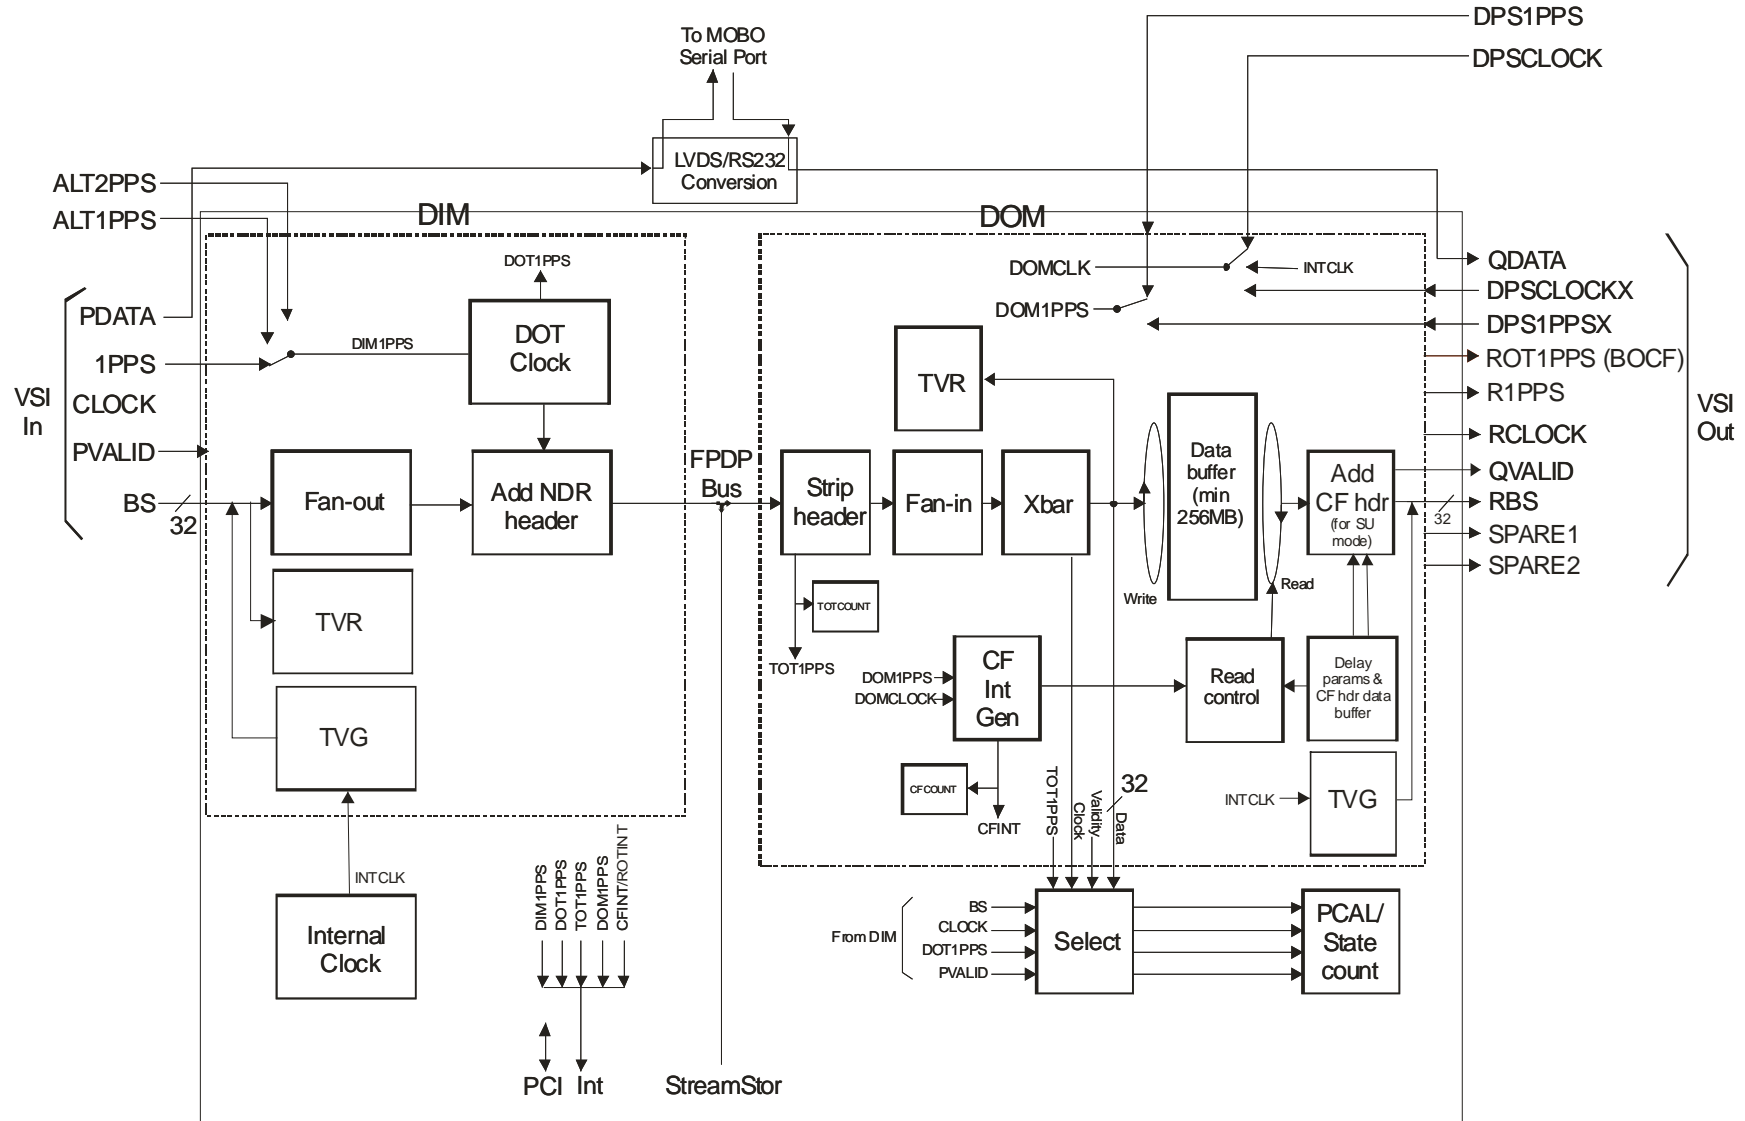


Figure 3: Strawman Mark 5B Block Diagram

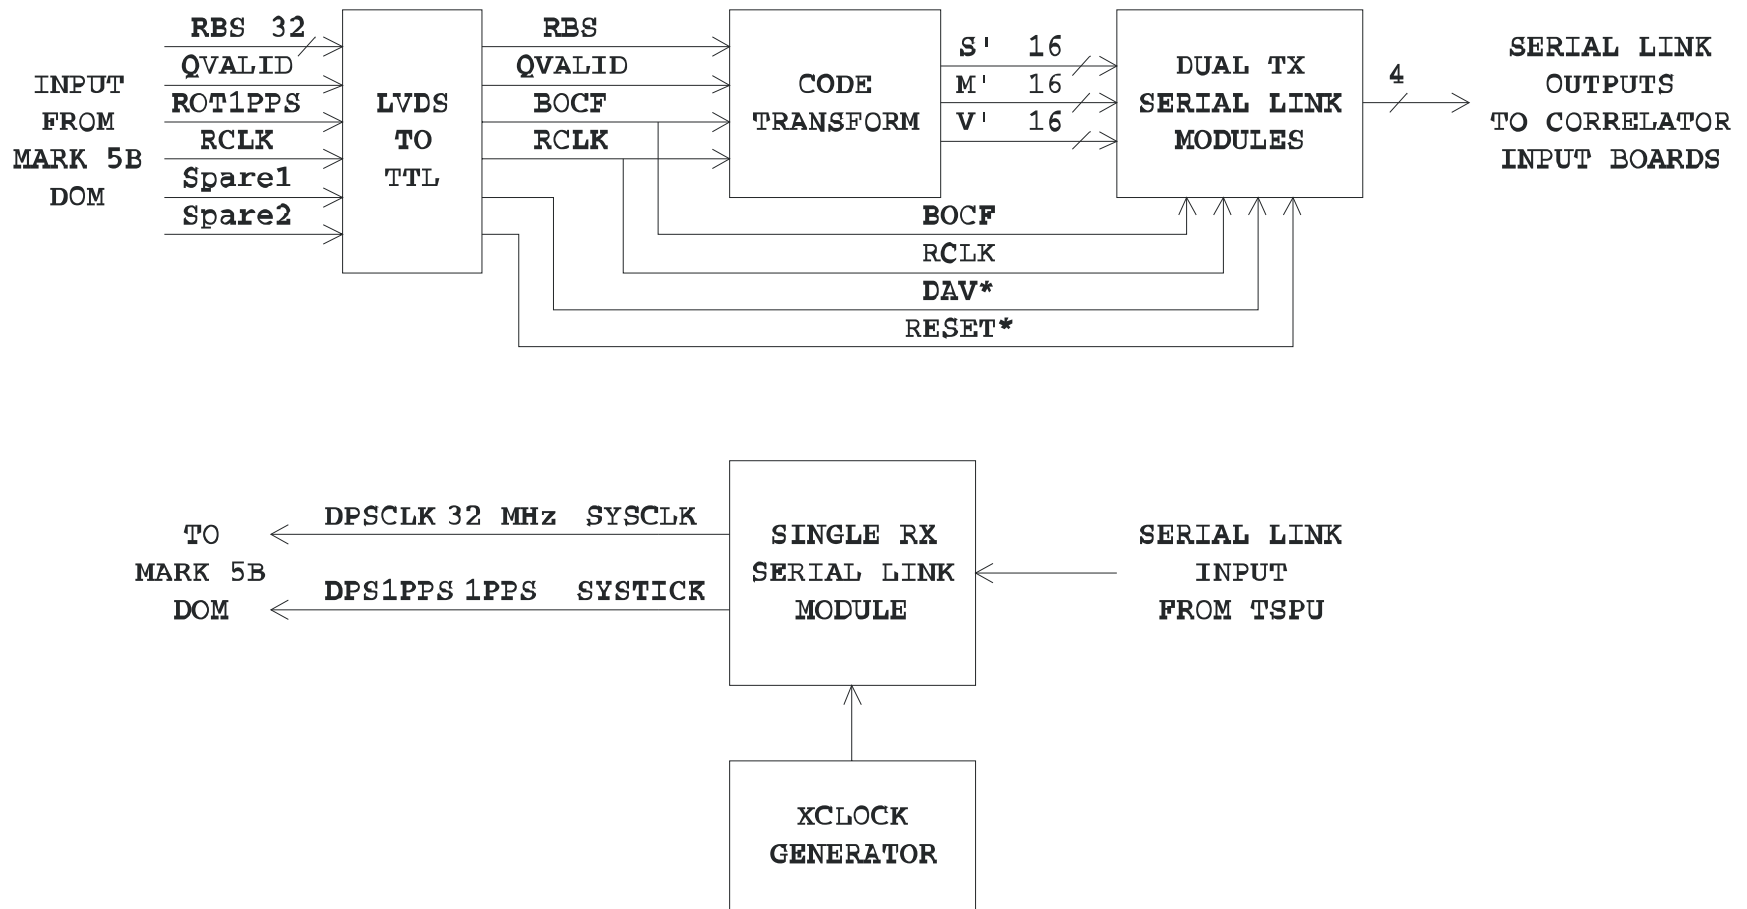


Figure 4: Correlator Interface Board (CIB) Block Diagram