ALMA Memo No. 298

Timing and Synchronization

Larry R. D'Addario
2000-Mar-09

Abstract

This report describes the current design of the timing distribution subsystem for the ALMA telescope, including facilities for synchronizing everything from the highest frequency local oscillators to the slowest switching cycles. Also considered is the need to maintain accurate knowledge of the relationship of the telescope's internal measure of time ("array time") to external measures like international atomic time.

1. REQUIREMENTS

Across the ALMA telescope, including the central building and the geographically dispersed antennas, synchronization among hardware devices is needed on a wide range of time scales for several different reasons. These include achieving coherence among local oscillators, establishing consistent switching cycles, ensuring consistent mode changes, and labeling the resulting measurements. Additionally, the long-term measure of time within the array must be accurately related to external measures of time; the main reason for this is to allow correct real-time evaluation of the ephemerides of astronomical objects of interest, but another reason is to allow coordination of ALMA's observations and their analysis with those of other observatories (the most stringent such requirement being the support of VLBI).

1.1 Internal Time and Synchronization

To achieve synchronization, it is necessary to distribute timing information. The distribution may be characterized by a few simple parameters: resolution, accuracy, and ambiguity interval. Resolution is the smallest time interval that can be distinguished in the distributed signals; (in)accuracy is the actual time difference between events that are nominally simultaneous (typically at different locations); and ambiguity interval is the longest periodicity in the distributed signals. (We say "distributed" because we envision a central master timing generator with all other devices being slaves, and this is the scheme planned for ALMA. But in principle a more symmetrical system is possible, where each device has an independent timing generator and the information is exchanged among them. For the most part, the same considerations would apply.) For example, consider using a sinusoid as the only distributed signal. The ambiguity interval is then the period of the sinusoid, one cycle being indistinguishable from another; the resolution is the smallest change of phase that can be measured, considering both the characteristics of the phase measuring instrument and any noise that accompanies the signal; and the accuracy is partly determined by the same phase measurement precision, but also by how well known is the transit time of the signal over the distribution path.

Accuracy may be specified in an relative sense, where fixed offsets are not counted as part of the error. In that case, it is primarily the stability of the timing (the extent to which the offsets are truely "fixed") that is of concern. This view is relevant when the offsets are determined and removed by offline astronomical calibration, so that they need not be zero in real time.

The need to maintain phase coherence among the LOs at the highest LO frequency sets the smallest resolution and (relative) accuracy requirement. For ALMA, these are both approximately 30 fsec (1/30 cycle at 1 THz). The need to synchronize switching cycles (phase, nutation, antenna position, etc.) among antennas and between antennas and correlator sets the largest ambiguity interval requirement. This is somewhat mode dependent, but is the order of 1 sec or less. If we take the minimum observational integrating time to be 1.0 msec, and let the maximum signal loss due to synchronization error be 0.1%, we obtain a requirement on absolute synchronization accuracy of 1 μsec.

1.2 External Time

The requirements for relating ALMA's time to external time are of two kinds: scale and initialization. The internal time scale must be such that ALMA's measures of time intervals and frequencies are in close

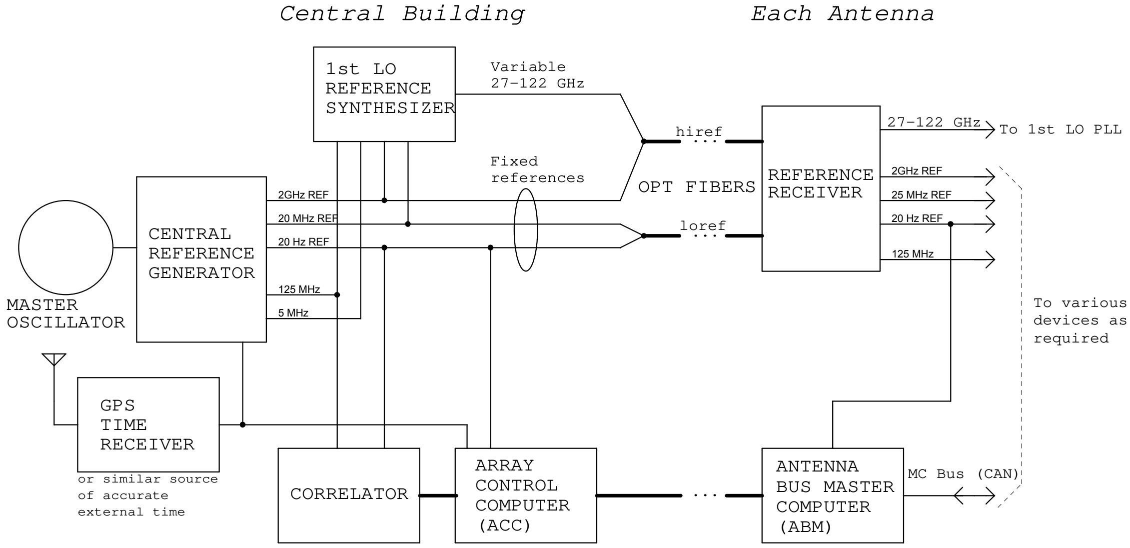

Figure 1: Simplified block diagram of the time distribution hardware.

agreement with international standards. Using hydrogen maser oscillators, agreement can be maintained within about $(10^{-12}\,\text{sec})/T$ over any interval T. By locking all timing signals in the array to one such oscillator, the same accuracy is maintained throughout. A long-term measure of time is then created by counting the cycles of the master oscillator. By careful initialization of this counter ("clock"), and by re-initialization when necessary, its reading can be made to track external time measures. For VLBI, knowledge of the difference between ALMA's time counter and standard time is expected to be accurate to ∼100 nsec; this can be achieved by comparison against GPS-distributed time. For all other known purposes, including time-labeling of observations, knowledge within 1 μsec should be sufficient.

2. DESIGN OVERVIEW

In the following discussion, numerical values from the current ALMA design will be used where appropriate. It is possible that some of these will be changed for practical reasons during development, but the underlying ideas should remain the same.

A simplified block diagram of the hardware involved in timing is given in Figure 1.

All timing in ALMA is coherent with a single master oscillator (probably a hydrogen maser) from which a hierarchy of timing reference signals is derived and distributed. From the smallest resolution up to an ambiguity interval of 50 msec, the distribution is via hardware that is part of the Local Oscillator (LO) subsystem. For longer ambiguity intervals, it is via hardware and software of the Monitor/Control (MC) subsystem and consists of delivery of commands to devices at the required times. The two methods are tied together because the hardware distribution includes a 20 Hz signal that marks the ends of each 50 msec interval, and because the MC subsystem can (by design) deliver any command within a known 50 msec interval. The 20 Hz reference and/or MC commands are available to any device that needs them. (Note that a "device" may be a computer.)

It is a major principle of this design that no two devices need to talk to each other in order to exchange timing information, even though they may need to be synchronized. Each receives timing signals from the master oscillator and master MC computer; nowhere else is timing information originated.

The hardware distribution is further divided into fixed-frequency and variable-frequency references. The only variable-frequency reference is for the first local oscillator of the receivers, and it is transmitted to each antenna at 27 GHz or above (to 122 GHz in the present baseline, and perhaps to 938 GHz in the future). This is the signal that needs 30 fsec stability. (It also has various other requirements that are beyond the scope of the present report.) The stability is achieved partly by including automatic compensation for changes in the path delay (on optical fiber) by monitoring the round-trip phase of an optical carrier (for details see ALMA Memo 267). The remaining references consist of several fixed-frequency signals.

The fixed-frequency references include 20 Hz, 25 MHz, and 2 GHz. These are chosen to facilitate unambiguous-phase synthesis of any other frequency that is a multiple of 20 Hz. This is equivalent to saying that they convey time with resolution $\ll$1/(2 GHz) and ambiguity interval <1/(20 Hz). In principle, transmitting only the 20 Hz and 2 GHz references is sufficient to achieve this; the intermediate 25 MHz is included because, in practice, it is difficult to transmit 20 Hz with sufficient stability and SNR to resolve one cycle of 2 GHz. At the antenna, these signals are used to synthesize the second local oscillator, fringe rotation offsets, digitizer clock, data transmission clock, and phase switching. They can also be used to control any other functions that require precise timing.

As mentioned, devices that require time with ambiguity longer than 50 msec must rely on delivery of commands by MC. Since the master array control computer (ACC) has access to a long-term counter related to external time, it can generate commands at known complete times. (Here "complete time" means a measure of time whose ambiguity interval is very large and can be ignored for all practical purposes — e.g., several centuries.) We call the long term counter the *master clock*, and we call its reading the *array time*. It is sufficient if the counter counts only the 20 Hz reference. Clearly a command to a device may contain data specifying the full array time on a specific 50 msec boundary, e.g. on the next one after receipt of the command. From then on, the device could count 20 Hz cycles and thereby implement a "slave clock" synchronized to the master. However, for robustness against power outages, maintenance activity, and other disruptions, the accuracy of slave clocks should not be relied upon over long intervals. Depending on the nature of the device, either the complete time should be sent by command whenever it is needed, or a command to re-set the slave clock should be sent periodically. In the latter case, the ambiguity interval is extended to the command period.

There are no devices in ALMA that are known to require slave clocks. This includes all computers other than the ACC. By avoiding slave clocks, considerable software simplification is achieved. Nevertheless, nothing in the design precludes their implementation.

Devices requiring high resolution timing to high accuracy, including all microwave synthesis (second LO, sampling clock, etc.), must use one of the higher frequency references, preferably the 2 GHz signal. However, for devices which require higher resolution than 50 msec but which do not need especially high accuracy, interpolation across the 50 msec periods is often sufficient. This is straightforward in most microcomputers: given a hardware interrupt on each 50 msec boundary, smaller intervals can be measured by internal timers using the processor's free-running (crystal) oscillator whose frequency is known sufficiently well. Much greater accuracy can be obtained by driving the processor from one of the distributed references (e.g., 25 MHz) rather than from a free-running oscillator.

In summary: (a) Time for <50 msec intervals is distributed via LO Subsystem hardware. (b) Time for >50 msec intervals is distributed via the MC Subsystem's communication hardware and software. (c) There is only one master oscillator and only one "clock" counter for long intervals; these determine array time. The relationship between array time and external time is determined at only one place.

3. DESIGN DETAILS

3.1 Signals Transmitted; Ambiguity Resolution Circuit

The first LO reference is transmitted as the difference between two CW laser signals on one optical fiber. This difference is phase locked to the master oscillator. One laser signal (the "master") is common to all antennas, and the other is separate for each subarray, so as to allow subarrays to operate on different frequencies. (See MMA Memo 267.)

The 2 GHz reference is transmitted by intensity modulation of the master laser signal of the first LO. It therefore benefits from the path delay stabilization circuitry and should have excellent phase stability. The lower-frequency references are transmitted separately, without path delay stabilization. However, they are "re-timed" upon receipt using the 2 GHz reference so as to inherit the latter's phase stability; details of this are explained in below and illustrated in Figure 2. Thus, at each antenna, all these signals have a fixed phase relationship to each other and to the master oscillator.

The 25 MHz reference must be transmitted with timing stability and resolution <500 psec so as to resolve one cycle of the 2 GHz reference. At each antenna, it is "captured" on a zero-crossing of the 2 GHz reference using a high-speed flip-flop. The captured signal has a period of 1/(25 MHz), but it

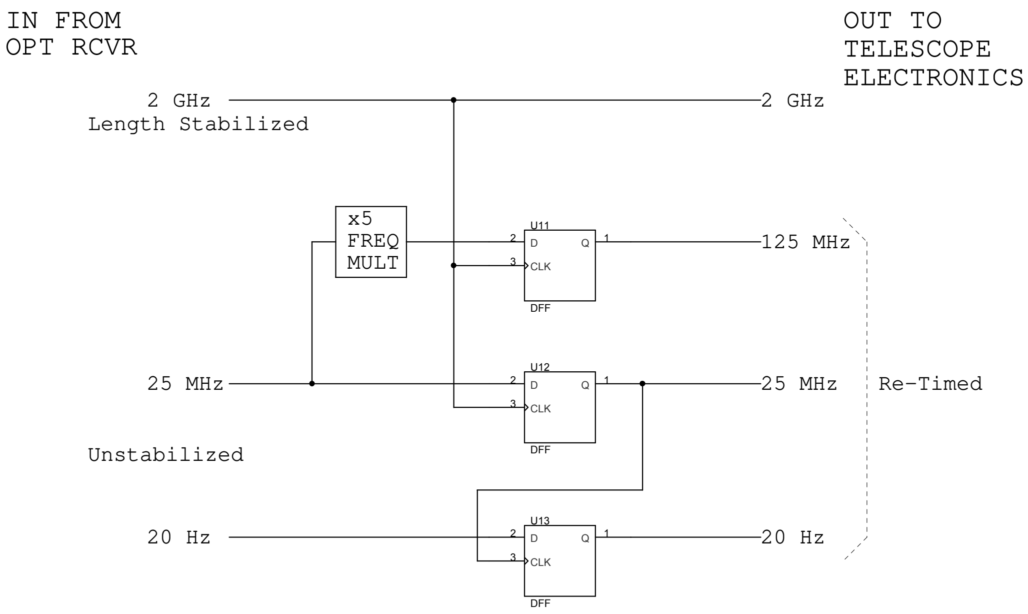

Figure 2: Ambiguity resolution circuit.

inherits the stability of the 2 GHz reference. The pre-capture 25 MHz signal may vary in phase by a large fraction of the 500 psec period of 2 GHz; this variation will be removed by the capturing process and will not be seen by subsequent circuits. By transmitting the 25 MHz reference over carefully installed fiber, the variation in propagation delay is expected to be $\ll$500 psec diurnally, and probably even seasonally, without special length stabilization hardware.

Similarly, the 20 Hz reference is captured on a zero-crossing of the captured 25 MHz reference. The 20 Hz signal received must then have an accuracy and stability $<1/(25\text{ MHz}) = 40$ nsec. This is easy with respect to delay variation, but it does mean that the 20 Hz signal must have good SNR. The captured 20 Hz signal will be distributed locally as a pulse (or square wave) whose rising edge defines a "timing event" that marks the boundary of a 50 msec interval.

3.2 Devices That Generate Non-Periodic Waveforms

Certain devices at the antennas deserve special discussion. These include the direct digital synthesizers (DDSs) that are part of the first LO and second LO PLLs, and that are responsible for fringe rotation and other fine phase control; devices for generating switching waveforms (including phase switching); and the sampling signal generator, whose phase must vary for fine tracking of the interferometer delay. (Much more detail about fringe rotation, phase switching, and DDSs is given in ALMA Memo No. 287.) In all these cases, the device must generate a waveform that varies with time in a complicated way that is not periodic. To accomplish this, we require that each such device operate as follows: Over a fixed interval of time equal to an integral number of 50 msec periods, the device generates the required waveform autonomously, to sufficient accuracy, by evaluating a fixed formula containing a small number of variable parameters. Prior to the timing event marking the beginning of each such interval, MC must deliver the parameter values to the device using one or more commands. During the 50 msec interval immediately preceding the beginning of an interval, MC must deliver a command indicating that the interval begins on the next timing event. The length of the interval depends on the requirements of the device and the way in which the waveform is parameterized, but it is expected that it will usually be $\sim$1 sec. Since these waveforms tend to depend on both observing mode and source location, there is no point in making the interval longer than the typical duration of an elementary observation.

3.3 Propagation Delay

The reference signals will suffer a significant propagation delay in transmission from the center to each antenna. This may be up to 120 μsec at the maximum distance ($\sim$25 km of fiber). The present design does not include any compensation for this delay, which means that absolute synchronization of the reference signals is not achieved among the antennas nor between an antenna and the center. For some purposes, this can be ignored. For others, including the timing of fringe rotation updates and (possibly) synchronizing of phase switching with the correlator, calculation of the required waveforms by the ACC must take account of the delay. A table of the actual delay to each antenna, based on measurements, will be maintained and available to the ACC. It is believed that knowing the delays to an accuracy of 1 μsec will be sufficient, and various methods can be used to do this. Re-measurement should be needed only when an antenna is moved or hardware is replaced, since long-term stability much better than 1 μsec is expected.

3.4 Time Tagging

When data is recorded in an archive (either astronomical data from the correlator or detectors, or technical monitor data, or logs of operational events), it is generally important that it be labeled with the (complete) array time. Whereas all archives will be physically written at the central building and not at antennas, and whereas all data from antennas will be transmitted in real time, it is recommended that all time tagging be done centrally. Computers responsible for such time tagging should have access to the master clock, either directly or through the ACC. Although slave clocks are not prohibited, their use is not recommended for the reasons discussed earlier.

3.5 Other Reference Frequencies

Some devices require periodic signals at frequencies different from those of the distributed references.

As mentioned, unambiguous-phase synthesis of any periodic signal whose frequency is a multiple of 20 Hz can be achieved from the distributed references. At relatively low frequencies and loose accuracy requirements, this can be done in the software of a microcontroller or real-time computer. There are a few cases where it will be done in dedicated hardware: (a) At the center, 5 MHz is needed by all first LO synthesizers so as to achieve 5 MHz tuning steps. (b) At the antennas, 125 MHz is needed for the first and second LOs, for the sampling clock generator, and possibly for other devices. (c) At the correlator, 62.5 Hz and 1000 Hz are needed to synchronize the internal timing with array time. In all these cases, the signal could be synthesized from the distributed 25 MHz and/or 20 Hz, but whereas it is needed by multiple devices the synthesis is done once at each location and distributed.