THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 54, No. 5, May–June 1975
Printed in U.S.A.

Digital Data System:

Network Synchronization

By B. R. SALTZBERG and H. M. ZYDNEY

(Manuscript received July 12, 1974)

The Digital Data System is synchronized by means of a network in the form of a master-slave tree. Functioning data links are used for distribution of synchronization signals and timing is recovered from the data stream. The master timing supply sets the frequency for the entire network. Each major node contains a nodal timing supply with an extremely long time constant, which phase-locks to the incoming data signal. The nodal timing supply contains frequency memory and sufficient accuracy to free run satisfactorily for several days if its inputs fail. A local or secondary timing supply is designed for operation in distant nodes of the tree. Each timing supply provides common clock signals to all DDS equipment in the office in which it is installed. All timing supply designs include a high degree of redundancy for reliability purposes.

I. INTRODUCTION

The network for the DDS consists of an interconnected set of digital transmission facilities. At any node of the network, if the average rate of transmission bits leaving a node is not exactly equal to the average rate entering the node, errors occur. Such errors are defined as slips, which means that bits are arbitrarily deleted if the input rate exceeds the output rate, or that bits are repeated or inserted arbitrarily if the input rate is slower than the output rate. The best performance is realized if no slips occur. This requires that every node in the system be synchronized to the identical average frequency and be capable of absorbing delay fluctuations.

The transmission media comprise radio, coaxial cable, and twisted pairs. Each of these admits to small but troublesome variation in delay because of thermal and other effects. The rates of change of such delay variations are equivalent to small frequency disturbances. To be free from slip under these conditions requires two elements: (*i*) reproduction of identical average frequencies at every node, and (*ii*) storage of sufficient data at each point of entry to accommodate the

inevitable delay variations of the transmission media. The details of elastic storage to accomplish the second function are discussed in the description of digital multiplexers.[1] The frequency distribution is the subject of this article.

II. SYNCHRONIZATION NETWORK

Distribution of accurate frequency for transmission systems has been the subject of many studies.[2] For DDS, a loosely coupled master-slave system was selected as the appropriate process to meet performance, cost, and administrative requirements. In normal operation, frequency information is transmitted down a topological tree which is a selected subset of the actual data facilities. Such a tree is shown in Fig. 1.

A reference frequency originates the timing signal at a location convenient for other Bell System purposes. This is directly transmitted to a master timing supply that creates a timing format compatible with the remainder of DDS. This signal is transmitted, via the 1.544-Mb/s DS-1 data stream, to major nodes containing nodal timing supplies. Each of these is phase-locked to the incoming data signal with a very long time constant so that the frequency generated at each node is extremely stable and always quite close to the average incoming frequency. The unit of phase controlled at each node is the 125-μs interval that represents a multiplexing frame. Elastic storage in each data-stream input withholds a sufficient number of bits so that data delivered to the node can be held in exact frame alignment

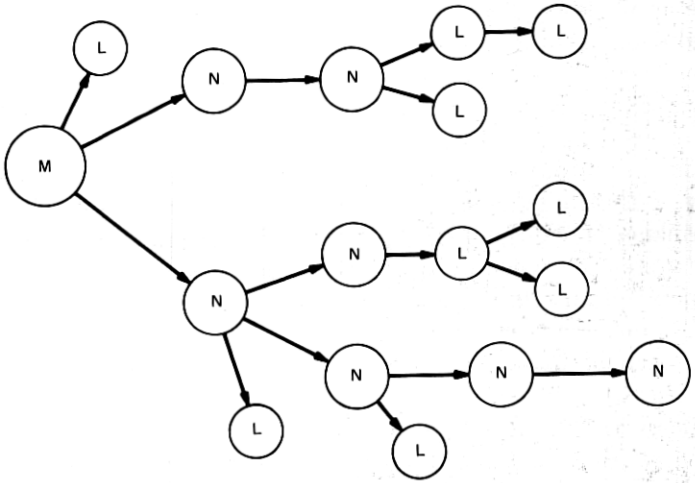

Fig. 1—Synchronization tree of master, nodal, and local timing supplies.

from each source, using the synchronization data stream as reference. This process is repeated at each major node in the tree. Minor nodes, with more limited connectivity, are equipped with phase-locked local timing supplies that are simpler in design and have somewhat more relaxed specifications, but which otherwise perform a similar function. When interconnected, such a system performs slip-free and is an efficient approach to timing distribution across the continental United States.

A primary concern in implementing a system of this sort is the possibility of a temporary interruption of transmission supplying a branch of the synchronization tree. Nodal timing supplies are designed to control the frequency of the internal oscillator by means of a digital proportional-plus-integral phase-locked loop in which the integral is held in a digital memory. This is described in detail in Section 3.1.2. Interface circuitry constantly monitors for the presence of a validly formatted timing signal in the data stream. If it is interrupted, further phase information is inhibited and the oscillator can then run free at the last remembered frequency stored in the integral memory. The interface circuitry automatically returns to normal operation when the outage is ended. If maintenance craftsmen determine that the outage will persist, a spare "hot" interface circuit and alternative transmission path is substituted. If a still more severe outage should occur, the nodal timing supply will run free of any timing input, beginning at its last frequency setting. During this condition, the timing supply relies on the precision of the oscillator to maintain synchronism. The selected oscillators hold frequency to an accuracy of one part in 10^{10} drift per day, which should result in a slip rate of less than one frame each day until the timing source is restored.

The local and secondary timing supplies are designed to operate from two independent, protected, 1.544-Mb/s channels, if they are available, and automatically transfer between signal sources. They are at a low level in the tree and a catastrophic failure causing loss of both input signals is unlikely to cause serious problems. They operate without digital storage of phase and contain only moderately accurate oscillators within their phase-locked loops. The secondary timing supply is physically arranged to permit easy conversion to a nodal timing supply when system growth so requires.

III. TIMING SUPPLY DESCRIPTION

3.1 Nodal timing supply

The nodal timing supply delivers a common set of clock signals to all DDS equipment in the hub office in which it is installed. It maintains synchronization with other hubs by phase locking to the framing

bit of an incoming DS-1 signal. When the incoming signal is interrupted or defective, the nodal timing supply free runs at the previous input frequency, appropriately averaged. The oscillator in the nodal timing supply phase-locked loop has sufficient long-term accuracy to permit satisfactory operation even after free running for several days.

A high degree of redundancy is provided for reliability purposes. As shown in Fig. 2, the interface unit, the phase-locked loop, and the output circuit are all duplicated. Units of each pair receive their power from different power supplies. Monitoring and control circuitry is provided to reconfigure the system in response to various outages or manual interventions. A suitable output will be present in the event of trouble.

3.1.1 Interface unit

The interface unit extracts the 8-kHz framing signal from an incoming DS-1 line signal. Each interface unit bridges on a T1 data multiplexer receive line. The two interface units are connected to the same line, to separate lines from the same origin, or to separate lines from different origins, depending upon the desired configuration of the synchronization tree. Only one interface unit is actually supplying a framing bit to the phase-locked loops at any given time. Manual switching between interface units is accomplished by means of a switch on a display-and-control panel. To avoid large phase hits when such a switch is made, a manual build-out switch is present on each interface unit for delaying the output in steps of one-twelfth of a cycle. The build out is adjusted during initial installation to bring the two interface unit outputs into as close alignment as possible.

In addition to the extracted framing bit, each interface unit also provides an output signal that indicates whether or not a valid framing bit is being extracted from the incoming line. A minor alarm is actuated if either interface unit is unable to extract a valid framing bit. The framing bit and the status indication from the selected interface unit are coupled to the phase-locked loops and to the control circuitry.

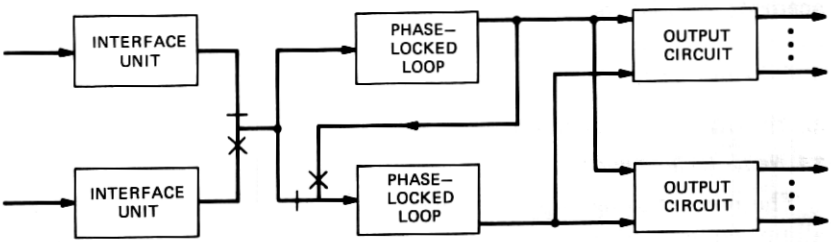

Fig. 2—Nodal timing supply block diagram. Only signal paths are shown.

3.1.2 Phase-locked loop

The phase-locked loop is a digital control system of the proportional-plus-integral type. The proportional-plus-integral phase-locked loop has the interesting property of producing no steady-state phase error in the presence of a steady frequency offset. More importantly, the phase-locked loop will remember its past operating frequency after the input to the loop is removed.

A block diagram of the phase-locked loop is shown in Fig. 3. The principal component is a 39A oscillator, which is a quartz oscillator mounted in a double oven. Its drift rate is less than one part in 10^{10} per day after one day of warm-up. A digital-to-analog converter is built into the oscillator to permit approximately linear fine control of the output frequency in response to a 14-bit parallel digital control input, weighted in a normal binary sequence. The 5.12-MHz output frequency of the oscillator may be varied over a total range of 0.8 part per million by means of this digital control. The least significant bit, therefore, changes the output frequency by a factor of approximately 5×10^{-11}.

The output of the oscillator is divided by a countdown chain to provide a large number of clock and gating signals with periods as long as 8.192 seconds. An 8-kHz and a 512-kHz output are fed to the output circuits. Other outputs from the countdown chain control the phase-locked loop serial arithmetic and provide clock signals to other parts of the timing supply.

In normal operation, both phase-locked loops lock to the output of the selected interface unit. The phase comparator compares the arrival time of this signal with a 4-kHz output from the countdown chain. The phase of alternate input pulses are measured to within 1/320 of an 8-kHz period. The measurement is formed by counting the number of pulses of a 2.56-MHz clock that occur between the start of the input pulse and a transition of the 4-kHz output. The phase comparator is, therefore, of the quantized sawtooth type.[3] The resultant phase measurement is read out serially during the other half of the

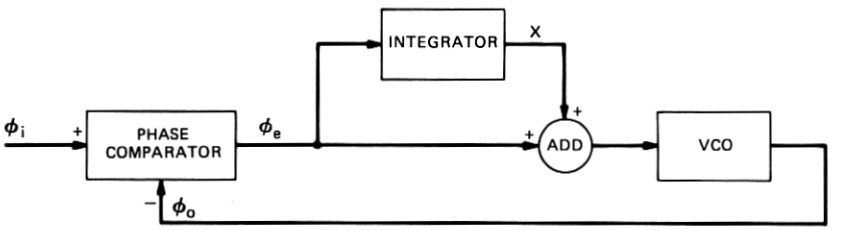

Fig. 3—Nodal timing supply phase-locked loop.

4-kHz signal, during which time no phase comparison is made. During each 8.192-second interval, the average of the 2^{15} phase measurements generated during the interval is calculated. Arithmetic is then performed on this average phase error during each interval to supply the control input of the oscillator. All arithmetic is in serial 2's complement binary form.

The arithmetic is performed in three steps during each 8.192-second interval. The average phase error is first transferred to temporary storage. Next, the average phase error is multiplied by 2^{-15} and added to the contents of an integral register. The integral register maintains a perpetual running sum of the phase error and may be manually reset to zero by means of a pushbutton on the faceplate of a circuit pack. The last step of the arithmetic consists of adding the average phase error, which was temporarily stored, to the contents of the integral register and storing the result back into a register. At the end of the operation, the 14 resultant bits are used to update the frequency control input of the oscillator. The oscillator input is updated only once every 8.192 seconds to reduce the effect on the lifetime of the relay switches in the built-in digital-to-analog converter.

Although the control loop is partially a sampled data system, the sampling interval is much shorter than any of the time constants involved. The summation in the integral register can be approximated very closely by an integral and the resulting open-loop gain function of the phase-locked loop may be approximated by

$$\frac{\Phi_o(s)}{\Phi_e(s)} = \frac{\alpha}{s}\left(1 + \frac{a}{s}\right).$$

The closed-loop gain function is, therefore,

$$\frac{\Phi_o(s)}{\Phi_i(s)} = \frac{\alpha(s + a)}{s^2 + \alpha s + \alpha a}.$$

The quantity α is determined strictly by the gain in the proportional path. In the normal mode described above, a phase error of 1/320 of a cycle at 8 kHz leads to a change of one least-significant bit at the oscillator control input. This, in turn, will change the oscillator output referenced to 8 kHz by a factor of 5×10^{-11}, or 0.4 μHz. α is, therefore, equal to 0.4 μHz divided by 1/320, or 1.28×10^{-4} second^{-1}. The quantity a is determined by the scaling factor of the input to the integral register multiplied by the frequency at which this updating occurs. Since, in the normal mode, the input to the register is multiplied by 2^{-15} and the updating occurs once every 8.192 seconds, a is equal to 3.73×10^{-6} second^{-1}.

The frequency response will show a very slight peaking effect so that some very low jitter frequencies may exhibit very slight amplification. Otherwise, since a is much less than α, the closed-loop response may be approximated by $\alpha/(s + \alpha)$. This is a simple single-pole response with corner frequency at 20.4 μHz for a one-sided noise bandwidth of 32.5 μHz.

The response to an input step in frequency Δf is given approximately by

$$\phi_e(t) \approx \frac{\Delta f}{\alpha} (e^{-at} - e^{-\alpha t}),$$

where the phase error is measured as a fraction of a cycle. The phase error, therefore, rises to approximately $\Delta f/\alpha$ with a time constant of $1/\alpha$, which is equal to 7810 seconds or 2.17 hours. The phase error then decays slowly to zero, with a time constant of $1/a$, which is equal to 269,000 seconds or 3.12 days.

After an interruption in the input signal, the phase-locked loop is forced to free run by setting the output of the phase comparator to zero. The output frequency of the oscillator in the free-running mode is then simply equal to α times the output of the integral register, which remains unchanged. While the input is present, the output of the integrator is given by

$$x(s) = \frac{a}{s} \Phi_e(s) = \frac{as}{s^2 + \alpha s + \alpha a} \Phi_i(s).$$

The corresponding frequency may therefore be approximated by

$$f_x(s) = \alpha x(s) \approx \frac{a}{s + a} \frac{\alpha}{s + \alpha} s\Phi_i(s).$$

This is equivalent to the tandem combination of two simple single-pole low-pass filters whose corner frequencies are at 0.593 μHz and 20.4 μHz. The first of these filters clearly dominates. One particular jitter frequency, which must be accounted for, is that due to daily variations in path length. Since this is a jitter frequency of 11.6 μHz, the daily jitter is reduced by a factor greater than 20 as far as the output of the integrator is concerned. When the input to the loop vanishes, the loop free runs at a frequency equal to the input frequency prior to the time of free running convolved by an exponential of time constant equal to 3.12 days.

Due to the extremely narrow filtering action described above and to previous line jitter, it is expected that the frequency offset immediately after free running should be less than five parts in 10^{11},

assuming that the frequency of the previous node is perfect. In addition, there is an initial frequency offset of up to five parts in 10^{11} owing to quantization effects. The initial frequency offset of a nodal timing supply immediately after it begins to free run, therefore, is less than one part in 10^{10}. In addition, the oscillator begins drifting at a rate of up to one part in 10^{10} per day. The nodal timing supply may, therefore, free run for up to 13.5 days before the frequency error is equal to the rate of one 8-kHz frame slipped per day. The accumulated phase drift will not reach $\frac{1}{2}$ cycle of an 8-kHz interval until 2.93 days have elapsed. Therefore, assuming zero initial phase error, the phase-locked loop will exhibit no slips if any outage in the input signal is restored before 2.93 days. However, whether or not any customer data will be slipped depends upon the initial fill of the T1 data multiplexer elastic store.[1]

The above parameters are unsuitable for initial lock-in of the loop during installation, or for restoring the loop after certain outages, owing to the narrow lock range and the long time constants involved. For this reason, a fast-start mode has been introduced which is manually entered by operating a switch on a faceplate of a circuit pack. Operation of the switch modifies the countdown chain in the phase-locked loop so as to change the gating waveforms used in the serial binary arithmetic. In the fast-start mode, the average phase error is multiplied by 2^5, thus increasing α to 4.1×10^{-3} second^{-1}. The input to the integral register is further increased by a factor of 2^9, thus changing a to 1.9×10^{-3} second^{-1}. The loop response is changed to a slightly under-damped one with a damping time constant of 488 seconds. Therefore, the loop may be brought into close lock in less than an hour in the presence of large phase or frequency errors.

3.1.3 Output circuit

Each of the output circuits accepts an 8-kHz and a 512-kHz waveform from each of the phase-locked loops and generates the composite waveform described in Section IV, Clock Distribution. The two output circuits are coupled so that at any given time both of them receive their inputs from the same phase-locked loop. When a defect is detected in the 8-kHz phase-locked loop output, the output circuits will switch their inputs to the other phase-locked loop and remain there until trouble is detected in that loop.

Sufficient fanout capability is provided and may be expanded to drive an arbitrarily large number of using bays. Each using bay receives its signal over two lines, one originating from each of the output circuits.

3.1.4 Monitoring, control, and alarms

Self-checking circuitry has been built into all major components of the timing supply. Any detected failure will generate at least a minor alarm. If trouble is detected in both phase-locked loops, in both output circuits, or in a combination of two or more power supplies that would render the system inoperative, than a major alarm is generated. An indicator light is provided to caution central office personnel that some manual control is in an abnormal condition, such as the fast start mode, a loop inhibited, or the manual free-run mode.

A built-in digital phase meter is included for manually checking the phase difference between any pair of 8-kHz waveforms in the timing supply.

Built into the nodal timing supply is a control algorithm that controls the input and output of the phase-locked loops in response to certain monitoring circuits. Associated with each phase-locked loop is a slip detector, which looks for an out-of-lock condition in the phase-locked loop. This is done by monitoring the phase difference between the phase-locked loop input and its 8-kHz output. If the phase difference proceeds from a region of moderate positive error to a region of moderate negative error, or vice versa, then the slip detector will generate a slip indication. This indication must be reset manually. In addition to the two slip detectors, a tracking detector is included which monitors the phase difference between the two phase-locked loop outputs and generates an output if the magnitude of this difference exceeds some small threshold. The fourth input to the control algorithm is the status signal from the interface unit in use. In the normal condition, both phase-locked loops receive their inputs from this interface unit. If the status indication goes off, the phase-locked loops are put into the free-running mode with loop A free running and loop B locked to the output of loop A. This keeps the phase difference between the two loop outputs small even though the system is in the free-running mode. If a slip in loop B is detected when in the free-running mode, appropriate indications are given and both loops will then free run separately since it is not possible to detect which loop is in error.

When in the normal mode, no action is taken when the tracking detector indicates a lack of tracking. However, when the no-track signal is on and a slip is detected, the loop that exhibited the slip is considered to be in trouble and its output is inhibited. The output circuits will automatically switch to the other loop if they are not already in that condition, and a minor alarm is given. If, after this event occurred, the other loop should also indicate a slip, a major

alarm is given. However, the output from the second loop is still supplied to the output circuits. Under no condition is the output of both phase-locked loops ever inhibited.

If, in the normal condition, a slip is detected while the tracking detector indicates the two loops are tracking, then the input signal is rejected even though the status indication from the interface unit is good. The timing supply is then put in the free-running mode with loop A free running and loop B locked to the output of loop A.

The nodal timing supply display-and-control panel includes light-emitting diodes imbedded in a block diagram of the system. These diodes indicate the states of various parts of the system.

3.2 Master timing supply

The master timing supply is identical to the nodal timing supply in all respects except for the interface units. The master timing supply is designed to lock on to a standard frequency transmission rather than the framing bit of a DS-1 signal.

A new Bell System Frequency Standard, using primary atomic standards sources, has been installed at Hillsboro, Missouri. The standard frequency transmission from Hillsboro to other offices in the country will be made via a 2.048-MHz pilot tone on L-multiplex facilities. The DDS master timing supply will be installed at St. Louis, which is the nearest DDS hub office to Hillsboro. The standard frequency will reach this office over two separate transmission facilities.

At the master timing supply, new special-purpose interface units are substituted for the normal interface units. These interface units count down the standard frequency to 8 kHz and keep track of the phase of this countdown process with self-checking. Normally, the output derived from the primary route is supplied to the phase-locked loops. If the primary route should have an outage, the input of the phase-locked loops is switched to the secondary route interface unit, in which the phase of the countdown has been previously adjusted to agree with the first interface unit. When the primary route is restored, the loop input is automatically returned to the first interface unit with phase continuity. Should both inputs to the master timing supply fail, the system will free run and the countdown circuits in the interface units will be reset so that upon restoration of either route the phase of the interface unit output agrees with that of the phase-locked-loop output.

3.3 Local (or secondary) timing supply

The local (or secondary) timing supply is a simple and low-cost system intended for use at those DDS local offices that have T1 data

multiplexers. It may also be used at small DDS hub offices, which are at the end of the synchronization tree. Its functions are identical to those of a nodal timing supply, the only difference being the use of a simpler phase-locked loop and control. The local timing supply is not able to free run without an excessive slip rate for more than a few seconds. As in the case of the nodal timing supply, fully redundant interface units, phase-locked loops, and output circuits are included. The interface units and the output circuits are identical to those in the nodal timing supply. However, since the local timing supply should not be permitted to free run for more than a few seconds, the interface units are arranged so that the input to the phase-locked loops automatically switches from one interface unit to the other when a failure in one interface unit or line signal is detected.

The fully redundant local timing supply is as reliable as the nodal timing supply. In very small local offices, where cost is the overriding consideration, it is possible to create a nonredundant local timing supply by installing only one interface unit, phase-locked loop, and output circuit. This arrangement, however, is only used where other nonredundant DDS elements are appropriate.

3.3.1 Phase-locked loop

Each local timing supply phase-locked loop includes a temperature-compensated 89A oscillator. This oscillator has a fractional frequency deviation of less than 12 parts per million under the expected extremes of aging and temperature. The 5.12-MHz output of the oscillator is divided down to provide the 8-kHz and 512-kHz outputs required by the output circuits, and to provide other waveforms required by the timing supply.

The phase-locked loop is a simple first-order type with a time constant of 1.04 seconds, so that its frequency response is that of a single-pole low-pass filter with corner frequency at 0.153 Hz. A sawtooth comparator generates a signal proportional to the phase error between the input signal from the interface unit in use and the 8-kHz output signal. The dc value of this signal is used to control the frequency of the oscillator. Only enough filtering is present in the loop to eliminate high-order modulation products from the phase comparator. This filtering has negligible effect on the loop response.

When the natural frequency of the oscillator is off by the maximum allowed value of 12 parts per million, the resultant static phase error of the loop is one-tenth of a cycle. If the input to the phase-locked loop fails, the loop will free run at its natural frequency and the phase will drift by the additional 0.4 cycle required for a cycle slip in 4.16 seconds. If the input is restored before this time, no slip will result.

Whether or not data slips occur depends on the fill of the T1 data multiplexer elastic store.

3.3.2 Monitoring, control, and alarms

As in the case of the nodal timing supply, the local timing supply also includes a display-and-control panel in which light-emitting diodes superimposed on a block diagram display the status of various components of the system. The same phase-metering circuit is also included. If a nonredundant system is installed, a feature built into the alarm logic changes the effect of some troubles that would normally produce a minor alarm to produce a major alarm. A permanent logic input is supplied by the second phase-locked-loop circuit pack. When this circuit pack is missing, the complementary logic level is supplied to the alarm logic. This changes the determination of minor or major alarm conditions and also causes those light-emitting diodes on the display panel that pertain to the second half of the system to be deactivated.

Associated with each phase-locked loop is a monitoring circuit. This monitoring circuit includes a slip detector which performs a function similar to that of the nodal timing supply slip detector and also requires a manual reset. In addition, other malfunctions in the phase-locked loop, in its input signal, or in the monitor itself are detected. In addition to the generation of alarms and indications, the phase-locked-loop monitor causes the phase-locked-loop output to be inhibited in the presence of certain trouble conditions. An end-of-range detector associated with each loop provides an indication when the phase error between the input and output of any loop is greater than one-quarter of a cycle in magnitude.

The only manual controls associated with the local timing supply are control keys that inhibit either of the phase-locked-loop outputs, and the slip-detector reset button. Interlock circuits on the phase-locked-loop outputs assure that under no condition arising from manual intervention or operation of the phase-locked-loop monitors can the output of both phase-locked loops be disabled simultaneously.

IV. CLOCK DISTRIBUTION

The cross-connection scheme for all DDS signals in an office requires that a common 64-kHz bit clock and a common 8-kHz byte clock be supplied to all equipment in the office. The waveforms required by the DDS equipment itself are shown in Fig. 4. The five-eighths duty cycle provides for a maximum tolerance to delays between pieces of DDS equipment and to difference in clock-propagation time from the timing supply to the different pieces of equipment.

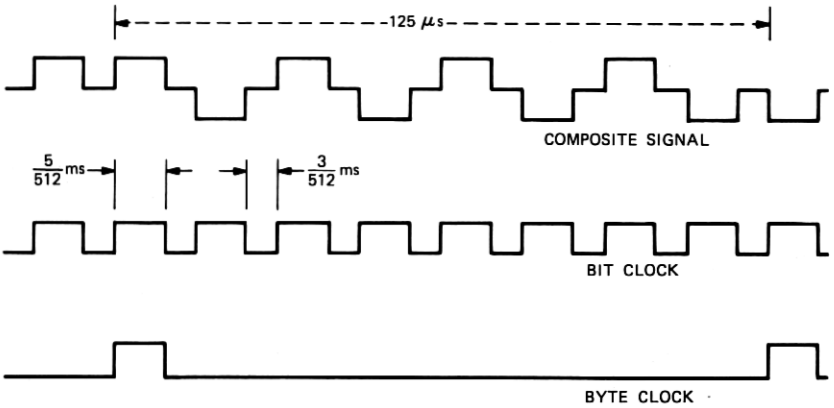

Fig. 4—Output clock waveforms.

The signal distributed by the timing supply to the using bays of DDS equipment contains information for both bit and byte clock combined in a single composite waveform. A composite waveform is transmitted redundantly over two balanced pairs to each using bay. The signal on each pair is a bipolar 64-kHz waveform having a five-eighths duty cycle. Each eighth pulse violates the bipolar rule in that it agrees in polarity with the preceding pulse. The basic waveform, therefore, provides the bit clock information, while the bipolar violation provides the byte clock information. Since the only information is carried in pulse presence and changes of pulse polarity, it is not necessary to keep track of the polarity of the two wires in the transmission line.

Each bay of DDS equipment contains a bay clock, power, and alarms unit, which serves several functions, one of which is to distribute the clock to all equipment in the bay. Two clock line terminators are included in this unit to receive the signals from the timing supply. Each terminator receives signals from both of the redundant pairs. The input composite waveform that is detected first is converted into the two unipolar clock signals required by the DDS equipment. Both terminators will continue to generate the clock signals even when one of the redundant distribution pairs or one of the timing supply output circuits is defective. The appropriate alarms and indications are given if either input signal is defective or if some defect in the circuit operation of the line terminator is detected. A switching circuit connects the outputs from one of the clock line terminators to a number of line drivers that distribute the bit and byte clock to all DDS equipment in the bay. Operation of the switching circuit depends on the presence of detected troubles.

The bit and byte clocks fed to all DDS equipment in an office are in phase, except for slight variations due to differences in propagation time from the timing supply to each bay.

V. ACKNOWLEDGMENTS

The initial synchronization plan was proposed by R. J. Deaton, J. F. Oberst, and R. B. Robrock. The nodal timing supply phase-locked loop was designed by G. Pasternack and J. M. Fallon, and the monitoring-and-alarm circuits by O. Napolitano. The local timing supply phase-locked loop and associated monitoring circuits were designed by R. C. Morris. The development of the interface unit and phase metering circuit was done by T. R. Lawrence and J. Kee. The output circuits and clock line terminators were designed by K. W. Boyd. Analysis of the synchronization network and its possible modes of failure was performed by N. H. Stochel. J. T. Ohlweiler performed the physical design for the nodal timing supply, and J. A. DeSena for for the local timing supply.

REFERENCES

1. P. Benowitz, S. J. Butterfield, M. P. Cichetti, Jr., and T. G. Cross, "Digital Data System: Digital Multiplexers," B.S.T.J., this issue, pp. 893–918.
2. J. W. Pan, "Synchronizing and Multiplexing in a Digital Communications Network," Proc. IEEE, *60*, No. 5 (May 1972), pp. 594–601.
3. C. J. Byrne, "Properties and Design of the Phase-Controlled Oscillator With a Sawtooth Comparator," B.S.T.J., *41*, No. 2 (March 1962), pp. 559–602.