IMPROVED TIME REFERENCE DISTRIBUTION FOR A SYNCHRONOUS DIGITAL COMMUNICATIONS NETWORK

Harris A. Stover
Defense Communications Engineering Center

ABSTRACT

A Time Reference Distribution Technique for synchronizing a digital communications network has been previously described. In that technique, time reference information from the highest ranking node in the network was supplied to all other major nodes over the best available path to each one, and provision was made for reorganizing the network to accommodate failures. By referencing the highest ranking node to a time standard such as UTC, that standard of time is made available at all major nodes with an accuracy dependent on the time transfer capabilities of the transmission paths. However, that technique fails to apply much of the information that could be used to provide a more accurate and stable (low phase fluctuation) system.

The improved Time Reference Distribution Technique described here also selects the highest ranking clock as master. In it, the natural hierarchy determined by the network connectivity is employed; instead of only supplying timing information over the best path to each node, the nodes determine the error in their local clocks either by using information from all neighboring (directly connected) nodes higher in the hierarchy than the local node or by using information from all neighboring nodes not lower in the hierarchy. To do this, information is exchanged between neighboring nodes and there are a set of procedures or rules for applying this information. These are presented and discussed.

Introduction

Digital communications have been growing at a very rapid rate and are being more widely applied. Switched digital communications networks in which the signals remain in digital form throughout the network are being developed. The problem of synchronizing such a network is much more complex than the synchronization of individual point-to-point digital communications systems. In digital point-to-point communications, it is only necessary that the receiver be correctly synchronized to the received bit stream and there is no need to synchronize the transmitters. In a network employing time division multiplexers and/or time division switches, each bit must be available at the multiplexer at the correct time to fill its assigned time slot in the interleaved bit stream. Since the bits to be interleaved at a time division multiplexer or switch originate at many locations throughout the network, it is important that their sources be adequately synchronized. Variable storage buffers can be placed in all received bit streams to act as reservoirs in which the bits are temporarily stored until they are needed. These buffers can accommodate variations in the transit time of the signal from one node to another and also small errors in the nodal clocks. However, it is preferable to use only a small portion of the available buffer capacity in normal operations and reserve most of its capacity for contingency situations. One suggested method of providing the desired transmitter synchronization is to distribute an accurate time reference through the network to all of its nodes [1, 2, 3, 4].

A switched digital communications network is normally made up of duplex (transmission in both directions simultaneously) digital transmission links which interconnect the nodes of the network. These duplex transmission links employ synchronization codes to allow the receivers to be easily synchronized to the received signals. The codes are chosen so as to be unlikely to occur as a part of the data sequence. This can be accomplished either by selection of unique patterns or by transmitting the synchronization code with a greater regularity than it would randomly occur. Since synchronization of the receiver for each individual transmission link to its received signal is required, such a network has a natural basis for the distribution of a time reference. Each node can use its local clock to control the time of transmission of the synchronization code from the local node. It can use the same local clock to measure the time of reception of the synchronization code from the other end of the transmission link. If the nodes at the two ends of the link exchange these measurements, they both can determine the difference between the two clocks with the effects of the signal transit time removed (except for usually small asymmetry in the two directions of transmission and instrumentation errors) [1, 2, 3, 4]. This is accomplished by simply subtracting one measurement from the other and dividing by two. This time comparison of clocks at neighboring nodes can be used to pass a time reference from the master to other nodes of the network.

A number of advantages for using an accurate time reference for synchronizing a digital communications network are given in references [1, 4].

Desirable Characteristics for a Digital Communications Timing Subsystem

Quite a large number of desirable characteristics can be listed for the timing subsystem of a large digital communications network [4]. Although it cannot be proved that these characteristics are necessary, many of them are widely accepted as desirable. These include: (1) Any node should be able to obtain all required timing information from its neighbors (directly connected nodes) without need to communicate with more distant nodes; (2) The timing subsystem should accommodate failures or destruction of major parts of the network and still remain operational; (3) Timing perturbations at a node should not propagate to other parts of the network (i.e., when one node makes a correction or other change in its clock, it is not desirable for this to propagate through the network like falling dominoes); (4) There should be no closed timing loop that could potentially contribute to system instability; (5) The timing subsystem should permit systematic self-monitoring to provide early detection of malfunctions so that they can be corrected before they interrupt communications traffic; (6) The communications timing subsystem should be compatible with other timing subsystems such as those employed for navigation; and (7) The timing subsystem should be self-organizing, initially and following failures.

A time reference distribution technique for a digital communications network as described earlier [1] provided these desirable characteristics, but did not consider one other important characteristic. An error in a clock at any level of the timing hierarchy should not affect the measurement of the clock error at any node lower in the hierarchy. Further, the selection of the paths for distribution of the time reference and the measurement of the error in the local clock should be independent of the correction of the errors in any of the other nodal clocks. This permits error correction in any clock to be made with minimum perturbations of the network while still not interfering with the accurate measurement of the error in any other clock. This can be accomplished very simply by having each node in the network inform its neighbors of the measured but uncorrected error in its own clock [2].

Distribution of the Time Reference Through the Network

Figure 1 will be used as an example digital communications network for the discussion of time reference distribution through the network.

The internodal connections of this example network, as shown by the straight lines, represent one third of the total possible internodal connections for a 12 node network. If any single node of such a network is selected as the master for the entire network, for timing purposes, the internodal links of the network cause the nodes of the network to fall into a natural hierarchy with the selected master at the highest level of the hierarchy. The second level of the hierarchy comprises all nodes connected to the master; the third level consists of all nodes connected to the second level but not to a level higher than second; etc. Figure 2 shows this hierarchy for the network of Figure 1 when node A is selected as the master, while Figure 3 shows the hierarchy when node E is chosen as the master.

In the time reference distribution technique described in earlier papers [1, 2, 3, 4], a simple set of rules was employed at each node. Some of these rules were used to determine the relative error in the clocks at neighboring nodes. Other rules were employed to establish the best paths to the master from each node. The rules used to establish the best paths were essentially the rules described by Darwin and Prim [5] for self-organizing master-slave timing systems. In order to apply these rules, each node is assigned a unique rank to be used in determining the order of succession to master and to help resolve ambiguities that could occur. Each transmission link is assigned a demerit value and each node chooses to receive its time reference through the particular neighbor that will provide the lowest demerit path between the local node and the master. When the numbers next to the individual links of Figures 2 and 3 are used to represent the demerit values for the individual transmission links, the best (lowest demerit) path from each node to the master is shown by the dashed lines in these figures.

It is obvious that for many nodes there are a large number of possible paths between that particular node and the master. In Figure 2, paths from the master to node I for which there is no backtracking through levels of the hierarchy include: ACI, ADI, AFI, AECI, AFCI, ABDI, ACFI, AEJI, AFJI, ACEJI, ACFJI, ABGJI, AFCEJI, AECFJI, and ADBGJI. By combining timing information passed over these various paths, it is possible to provide greater timing accuracy. Also, the network does not have to be reorganized following some types of failure that would require reorganization if only the best path to each node were used for the time reference distribution. However, for effective use of this information from many different paths a set of rules or procedures is needed.

Each transmission link used for a time reference transfer within the network will introduce some error in the comparison of clocks at the two ends of the link. For a large number of links of a given type, these errors in the comparison of the clocks can be assumed to be random with a mean value of zero. Therefore, the inaccuracy of a link can be characterized by the standard deviation (or variance) of the expected error. (Inaccuracy as used here refers to the inaccuracy, or error, of a clock error measurement.) Figure 4 illustrates a tandem connection of five links. Since the error (in the measured difference between two clocks) associated with each of these links (in other parts of the text these errors are referred to as inaccuracies of clock error measurements) is assumed to come from a random distribution with zero mean, the errors (inaccuracies) statistically add as the square root of the sum of the squares. Therefore the error for the tandem connection can be characterized as equation (1), where the E's represent the standard deviations and E^2 is a variance.

$$E_{AF} = \sqrt{E_{AB}^2 + E_{BC}^2 + E_{CD}^2 + E_{DE}^2 + E_{EF}^2} \qquad (1)$$

Figure 5 illustrates two nodes connected by two paths in parallel. In this case, we would expect a combined accuracy that would be statistically better (based on more samples) than that of either path by itself. Let the measurement made over the first of the two paths be $M_1 = V + E_1$, where V is the true value and E_1 is the error introduced by the first path. Similarly, let the measurement over the second path be $M_2 = V + E_2$. It is desirable to weight and combine these two measurements in such a way as to obtain the statistically most accurate measurement. See equation (2).

$$M_{AB} = W_1 M_1 + W_2 M_2 = W_1 (V+E_1) + W_2 (V+E_2) \qquad (2)$$

The weighting factors W_1 and W_2 apply to both the true values and the errors. Although the weighted true values add linearly, the weighted error values (being random) add as the square root of the sum of the squares, so that:

$$M_{AB} = (W_1 + W_2) V + \sqrt{W_1^2 E_1^2 + W_2^2 E_2^2} \qquad (3)$$

Since it is desired that the combined result M_{AB} be the true value with a statistically minimum error, W_1 plus W_2 must be equal to 1, and the expression under the radical sign must be made minimum by the selection of W_1 and W_2. Substituting $(1-W_1)$ for W_2 in (3) and finding the value of W_1 that minimizes the statistical error gives (4), and subtracting this value of W_1 from 1 gives W_2 as shown in (5).

$$W_1 = \frac{E_2^2}{E_1^2 + E_2^2} \tag{4}$$

$$W_2 = \frac{E_1^2}{E_1^2 + E_2^2} \tag{5}$$

Putting these weighting factors in (2) gives (6):

$$M_{AB} = \frac{\dfrac{M_1}{E_1^2} + \dfrac{M_2}{E_2^2}}{\dfrac{1}{E_1^2} + \dfrac{1}{E_2^2}} \tag{6}$$

By using this combined value as given by (6) as one member of a new parallel pair, computing a new combined value, and repeating this procedure until the total number of parallel paths are included, the desired weighting factor for the measurement from the path p of n parallel paths can be written as (7).

$$W_p = \frac{\dfrac{1}{E_p^2}}{\displaystyle\sum_{i=1}^{n} \frac{1}{E_i^2}} \tag{7}$$

The resulting statistical error for the combined measurement based on n parallel paths as found by substituting the weighting factors of (7) into (3) is given by (8).

$$E_c = \frac{1}{\sqrt{\displaystyle\sum_{i=1}^{n} \frac{1}{E_i^2}}} \tag{8}$$

The information of equations (1) through (8) can be applied in combining timing information passed over several different paths. This will improve the timing accuracy at some nodes of the network, but the application of the information should follow a suitable set of rules.

Rules for Time Reference Distribution Via Multiple Paths

Some basic considerations that influence the choice of timing information to be transferred within the network and the selection of a particular set of rules for using this information will be discussed prior to the presentation of a suitable set of rules.

Rather than transferring all timing information to a single location where a common processor can be used for all timing information from all nodes, it is much simpler and more reliable if each node of the network can receive all required timing information from its neighbors (directly connected nodes) and use this information in a rather simple local processor (microprocessor). In a network employing a centralized processor, the common processor becomes a point of high vulnerability that reduces network reliability and survivability; and the efficiency of utilization of transmission facilities is reduced because of the large amount of information that must be transferred to and from the central processor in order to serve nodes throughout the network. Therefore, it is desirable that all timing information required by any node either be stored at that node or be supplied by its neighbors without any need for any node to communicate with nodes more distant than its own neighbors.

Consider combining timing information at node B of Figure 6 that comes over the paths AB and ACB. If this combined measurement at node B is then used to determine a combined timing measurement at node C, this new node C measurement could then be used to determine a new combined measurement at node B which could be used to determine a new one at C, etc. The resulting iterative process would change the timing at nodes B and C over a large number of iterations without introducing any new measurements from node A. The passing of information back and forth between nodes B and C cannot improve its accuracy, but could possibly introduce additional error due to the link BC each time the link is traversed. It is desirable to provide rules that will make effective use of combined timing information from multiple paths while preventing such iterations. These iterations can be avoided if each node is prevented from using timing information that has been previously influenced by that same node. To accomplish this and still make effective use of timing information over many multiple paths, two classes of timing information can be maintained at each node. Class 1 timing information (clock error measurements and inaccuracy values for those measurements) is based only on information received from nodes higher in the hierarchy than the local node, while Class 2 timing information is based on information received from all nodes not lower in the hierarchy (those at the same level in addition to those higher). If only Class 1 information is used, there will be no closed feedback paths, but useful information from other nodes at the same level in the hierarchy will not be used in determining the clock errors. If, when deriving Class 2 timing information a node uses Class 2 information from other nodes higher in the hierarchy but is restricted to only Class 1 information from nodes at the same level in the hierarchy, the closed paths will still be avoided and the undesirable iterations will be avoided. Some of the other considerations in selecting the particular system were listed in a previous section.

In order to provide an effective self-organizing method of accurately distributing a time reference through a digital communications network, information is exchanged between neighboring nodes. This information is applied in compliance with a set of rules. These rules are discussed in the text after the following list of information which is transmitted by each node to its neighbors:

INFO 1. Rank of the clock used as the master time reference for the local clock. (This information is used to assure that the highest ranking clock in the network is used as master and to establish the order of precedence to master when a master fails.)

INFO 2. Number of links between the local node and its master time reference. (This information is used to establish the desired hierarchy.) ("Local node" is used in this discussion as a particular node under discussion which transmits information to its neighbors and receives information from them.)

INFO 3. Time of the clock at the remote end of the link (including the effect of the time required for the signal to transit from the remote node to the local node) as measured by the clock at the local node. (This information is used to determine the difference between the clocks at the two ends of the link with the signal transit time removed from the comparison.)

INFO 4A. Measured but uncorrected error in the local clock based on information from those neighbors higher in the timing hierarchy than the local node. (The term measured error as used here includes errors obtained by mathematically combining other measurements and the error in this measured error will be called its inaccuracy.) (The resulting Class 1 error information is passed to all neighbors not lower in the timing hierarchy than the local node and can be used by neighbors at the same level to determine their Class 2 errors, or by nodes at higher levels to aid in system monitoring.)

INFO 4B. Same as 4A, except based on information from those neighbors not lower in the hierarchy than the local node (Class 1 information from the same level and Class 2 from higher levels). (The resulting Class 2 error information is passed to all neighbors lower in the hierarchy than the local node for use in determining their Class 1 errors.)

INFO 5A. Estimated inaccuracy, stated as a variance (or standard deviation), of the local clock based on information from all neighboring nodes higher in the hierarchy than the local node. (The resulting Class 1 information is passed to all neighbors not lower in the timing hierarchy than the local node, and it can be used by neighbors at the same level to determine weighting factors (see discussion of equation 7) for combining INFO 4B information from their neighbors into a Class 2 measured error.)

INFO 5B. Same as 5A except based on information from all neighbors not lower in the timing hierarchy than the local node. (This resulting Class 2 information is passed to all neighbors lower in the timing hierarchy than the local node for use in determining weighting factors (see discussion of equation 7) for combining INFO 4A information into a Class 1 measured error.)

Notice that the two classes of information under item 4 and the two classes of information under item 5 are distinguished by the sources of information used to obtain them and also by the nodes that make use of them. When the rules for their use are also considered, it will be observed that they prevent the formation of the closed feedback paths.

Each node applies the following set of rules for the use of the information received from its neighbors.

Rule 1 A node initially entering the network will use its own clock as its time reference until a better reference can be determined. Its own clock provides a basic time reference to which the node always returns when it has no better reference available. Under these conditions, the local node supplies the rank of its own clock to its neighbors as INFO 1.

RULE 2 The first type of information received from neighboring nodes, INFO 1, provides the local node with the rank of the clock used as the master time reference by each of its neighbors. If one or more neighbors' reference clocks outrank the local clock, the node will select those neighbors (or the single neighbor) referencing the highest ranking master and use them (or it) in determining its own time reference, i.e., measuring the error in its own clock. The rank of the master time reference used by the selected neighbors will be supplied to all neighbors as INFO 1, i.e., the rank of the clock used as master for the local node. Continued application of this rule by all nodes will result in all nodes referencing the same highest ranking master clock.

RULE 3 If the local node is referencing its own clock there are no links between the local node and its master reference and this information is supplied to its neighbors as INFO 2. The second type of information, INFO 2, as received from its neighbors provides the local node with information about the number of links between each neighboring node and that neighbor's master time reference. Unless the clock at the local node outranks the master reference of all of its neighbors, the number of links between the local node and the master is greater by one than that of the neighbors (or neighbor) selected by rule 2 which have the least number of links between themselves and their master. This information is supplied to the neighboring nodes as INFO 2. Continued application of this rule will result in establishing the desired natural hierarchy such as shown in Figures 2 and 3. INFO 2 information as transmitted to neighboring nodes and as received from them indicates the position in the hierarchy of the local node relative to each of its neighbors.

RULE 4 The third type of information, INFO 3, as received from neighboring nodes provides the local node with the time difference between the local clock and the clock at each neighboring node (including the signal transit time from the local node to the neighboring node.) INFO 3 as transmitted to the corresponding neighboring node is subtracted from this information, and the difference is divided by 2. This provides a measurement of the actual time difference (no transit time included) between the local clock and the clock at each neighboring node [1, 2, 3, 4].

RULE 5 Each neighboring node not higher in the hierarchy than the local node transmits to the local node, as INFO 4A, the Class 1 measured but uncorrected error of its own clock, i.e., the error determined using information from that neighbor's neighbors that are higher in the hierarchy than the neighbor. Similarly, each neighbor higher in the hierarchy than the local node transmits to the local node, as INFO 4B, the Class 2 measured but uncorrected error in its own clock, i.e., the error determined using information from that neighbor's neighbors that are not lower in the hierarchy than the neighbor. As received, this information gives a measured but uncorrected error for each neighboring node. To this is added the difference between the local clock and each neighboring clock as determined by rule 4. The result is a set of error measurements for the local clock based on information from each of its neighbors. (The reason for using Class 1 information from some neighbors and Class 2 information from others is to avoid closed feedback paths while still making very effective use of the available information.)

Rule 6 Each neighboring node not higher in the hierarchy than the local node transmits to the local node, as INFO 5A, the estimated inaccuracy, stated as a variance (or standard deviation), of its Class 1 measured error. Similarly, each neighboring node higher in the hierarchy than the local node transmits to the local node, as INFO 5B, the estimated inaccuracy, stated as a variance (or standard deviation), of its Class 2 measured error. This information, as received, is the estimated inaccuracy of the measured but uncorrected error associated with each neighboring node. Add to each member of this set of information (directly if stated as variances or as the square root of the sum of the squares if stated as standard deviations) the estimated inaccuracy of the link between each neighbor and the local node as determined during engineering design. The result is a set of inaccuracies for the set of measured errors in the local clock based on information from each neighbor.

The estimated inaccuracy attributed to the link between the local node and a neighbor as established during engineering design includes several parameters. It includes an effect due to the differences in signal transit time in the two directions of the duplex link which includes delay differences in the transmitters and receivers at the two ends of the link. It also includes inaccuracies in the equipment used to measure timing differences between the received signal and the local clock.

Rule 7 From the set of error measurements for the local clock as determined by rule 5, and the associated inaccuracies determined by rule 6, only those for neighbors higher in the timing hierarchy than the local node are selected. These error measurements are combined according to equation (6) to determine a Class 1 measurement of the error in the local clock, i.e., one based on neighbors higher in the hierarchy than the local node. This is supplied as INFO 4A, the measured error in the local clock, to all neighbors <u>not</u> lower in the timing hierarchy than the local clock.

<u>Rule 8</u> From the set of inaccuracies determined by rule 6 only those for neighboring nodes higher in the hierarchy than the local node are selected. These are combined according to equation (7) to determine the inaccuracy for the measured error in the local clock based on information from neighbors higher in the hierarchy than the local node. This information is supplied as INFO 5A to all neighbors <u>not</u> lower in the timing hierarchy than the local node.

<u>Rule 9</u> From the set of error measurements for the local clock as determined by rule 5 and the associated inaccuracies determined from rule 6, all those for neighbors <u>not</u> lower in the hierarchy than the local node are selected. These error measurements are combined according to equation (6) to determine a Class 2 measurement of the error in the local clock, i.e., one based on all those neighbors <u>not</u> lower in the timing hierarchy than the local node. This is supplied as INFO 4B, the measured error in the local clock, to all neighbors lower in the timing hierarchy than the local clock.

<u>Rule 10</u> From the set of inaccuracies determined by rule 6 all those for neighbors <u>not</u> lower in the timing hierarchy than the local node are selected. These inaccuracies are combined according to equation (7) to determine the inaccuracy for the measured error in the local clock based on information from all those neighbors <u>not</u> lower in the timing hierarchy than the local node. This inaccuracy information is provided as INFO 5B to all neighbors lower in the timing hierarchy than the local node.

The combining of information over several different paths, in addition to providing more accurate time measurements at many nodes remote from the master, reducing the need for massive reorganization of the network following some failures as required when using only the best path, also provides the possibility for quantitative evaluation of the fitness of the timing subsystem. Since each time error measurement (the term measurement as used here includes the mathematical combination of measurement information from different sources) has a corresponding estimated inaccuracy, these time error measurements and their corresponding inaccuracy estimates can be used to provide a quantitative alarm system. This leads to rule 11.

<u>Rule 11</u> Rule 5 provides a set of error measurements for the local clock based on information from each neighboring node. Rule 6 provides a corresponding set of inaccuracies for these error measurements. Rule 9 provides a combined measurement for the error in the local clock. Rule 10 provides a corresponding inaccuracy for the

combined measurement. The combined measurement as determined by rule 9 is subtracted from each member of the set of error measurements determined by rule 5. The resulting set gives the difference between each individual measurement and the combined measurement. The inaccuracy (stated as a variance) determined by rule 10 is added to each member of the set of inaccuracies obtained by rule 6 (also stated as a variance) and the square root of each member of this set is taken to obtain a set of estimates of the standard deviations of the clock error measurements based on information from each neighbor relative to the combined clock error measurement. Each member of the set of differences between individual measurements and the combined measurement is divided by the estimate of the corresponding standard deviation to obtain a normalized set of ratios. The lowest level alarm could be activated when the ratio reaches 2. This would not be very significant because this ratio would have approximately a 5% probability of occurrence in a normally operating system. A second level alarm when the ratio reaches 3 should be quite significant since its probability of occurrence in a normally operating system should be only about 0.3%. A third level alarm when the ratio reaches 4 should initiate some form of a problem investigation since its probability in a normally operating system should be less than 0.01%. A fourth level alarm when the ratio reaches 5 should initiate definite corrective action since its probability in a normally operating system might be expected to be less than one in a million.

There are also other capabilities for checking for an erroneous information exchange. For example, the node serving as master should reference its own clock and inform its neighbors that there are zero nodes between itself and its master. Every node connected directly to the master should inform its neighbors that it has one link between itself and its master. If any neighboring node tells the master that there is other than one link between itself and the master, the master can interpret this as detection of a problem. For every node in a stabilized operating system, either each neighboring node should be reporting the same number of links between itself and the master as the local node, or it should be reporting one more or one less than the local node. Any node receiving a report from one of its neighbors that the neighbor's distance from the master differs by more than one from the local node's distance from the master has detected a problem that should either cause reorganization of the network or other corrective action.

The procedure presented here for providing a very good time reference distribution through a digital communications network is based on an assumption that the different paths passing the time reference between two nodes are independent, i.e., do not share any of the same transmission links. Although this independence does not always exist, the degradation due to the dependencies that do exist should generally be acceptable. It is this assumption of independence that makes it possible to provide an accurate time reference distribution through the network by the application of simple procedures and calculations at each node using only simple information from neighboring nodes. In order to allow fully for dependent paths through the network, it would be necessary to keep a record of all dependent paths and make the necessary information available at every location where a calculation involving a particular dependent path is made. This would impose a very great increase in both communications and computation support to the timing subsystem. Because of the extensive increase in computation and communications required to permit full consideration of dependent paths, it is recommended that the independence assumption be made and that the simple procedures described above be applied to obtain improvements over time reference distribution via only the best path. The assumption of independent paths will indicate an apparent accuracy which is somewhat greater than the actual accuracy. The actual accuracy will usually lie somewhere between the accuracy obtained using only the best path and the accuracy indicated by using the above rules. One method of partially compensating for the effect of dependent paths might be to select a typical network arrangement using typical link inaccuracies, calculate the inaccuracy at each node by the above rules and also calculate it taking dependent paths into consideration. The average difference between the two methods could be determined for each level in the hierarchy and stored at every node. Then this average value for the local node's level in the hierarchy could be added to the value obtained using the set of rules above. The result might be expected to be statistically more accurate than just accepting the value using the assumption of path independence with no attempt to compensate.

Table I shows the inaccuracies, expressed as variances, for the measurement of the local clock errors for each node of Figure 1 when node A is the master as shown in Figure 2. Table II shows the inaccuracies when node E is the master as shown in Figure 3. These evaluations were obtained using the link inaccuracies (expressed as variances) shown in Figures 2 and 3 by applying the rules given above. The application of the rules requires that the Class 1 inaccuracy at a node is obtained from the Class 2 inaccuracies of nodes higher in the hierarchy than the local node while the Class 2 inaccuracy at a node is obtained by combining the Class 2 inaccuracies of neighboring nodes higher in the hierarchy than the local node with the Class 1 inaccuracies of neighbors at the same level in the hierarchy as the local node. In each case the local clock error measurement with Class 2 inaccuracy is available for use at the local node. As observed from the tables, the local clock error measurement with Class 2 inaccuracy is nearly always more accurate than that obtained using the best path. Because dependent paths were not taken into consideration the actual inaccuracy is probably between these two.

TABLE I. ERROR ESTIMATES FOR NETWORK OF FIGURE 2

Node	Error Estimate for Time Reference Via Best Path	Class 1 Error Estimate	Class 2 Error Estimate
A	0	0.000	0.000
B	2	2.000	1.333
C	1	1.000	0.652
D	2	2.000	1.333
E	1	1.000	0.833
F	1	1.000	0.750
G	3	1.645	1.243
H	2	1.833	1.833
I	3	0.931	0.647
J	2	1.082	0.618
K	3	1.699	1.123
L	4	3.243	1.473

TABLE II. ERROR ESTIMATES FOR NETWORK OF FIGURE 3

Node	Error Estimate for Time Reference Via Best Path	Class 1 Error Estimate	Class 2 Error Estimate
A	1	1.000	0.833
B	3	2.833	1.326
C	2	4.000	1.333
D	3	2.833	1.086
E	0	0.000	0.000
F	2	0.848	0.693
G	4	2.143	1.193
H	1	1.000	1.000
I	3	1.773	0.850
J	2	2.000	2.000
K	2	2.000	1.440
L	3	1.383	1.383

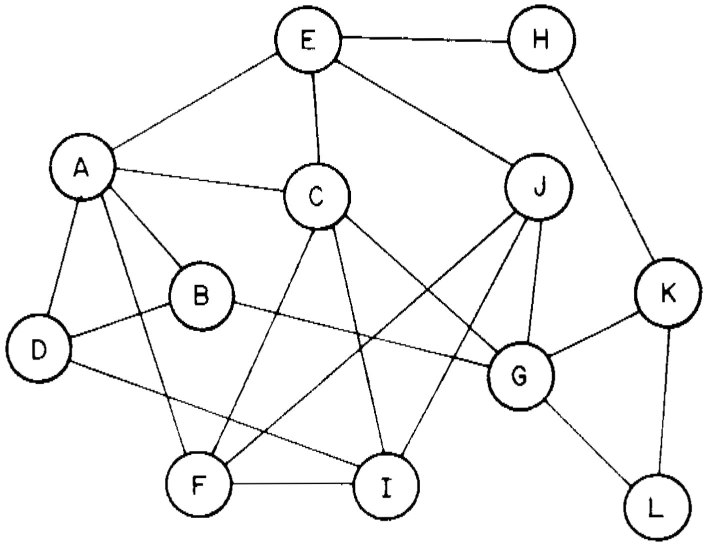

FIGURE 1. EXAMPLE OF A 12 NODE NETWORK

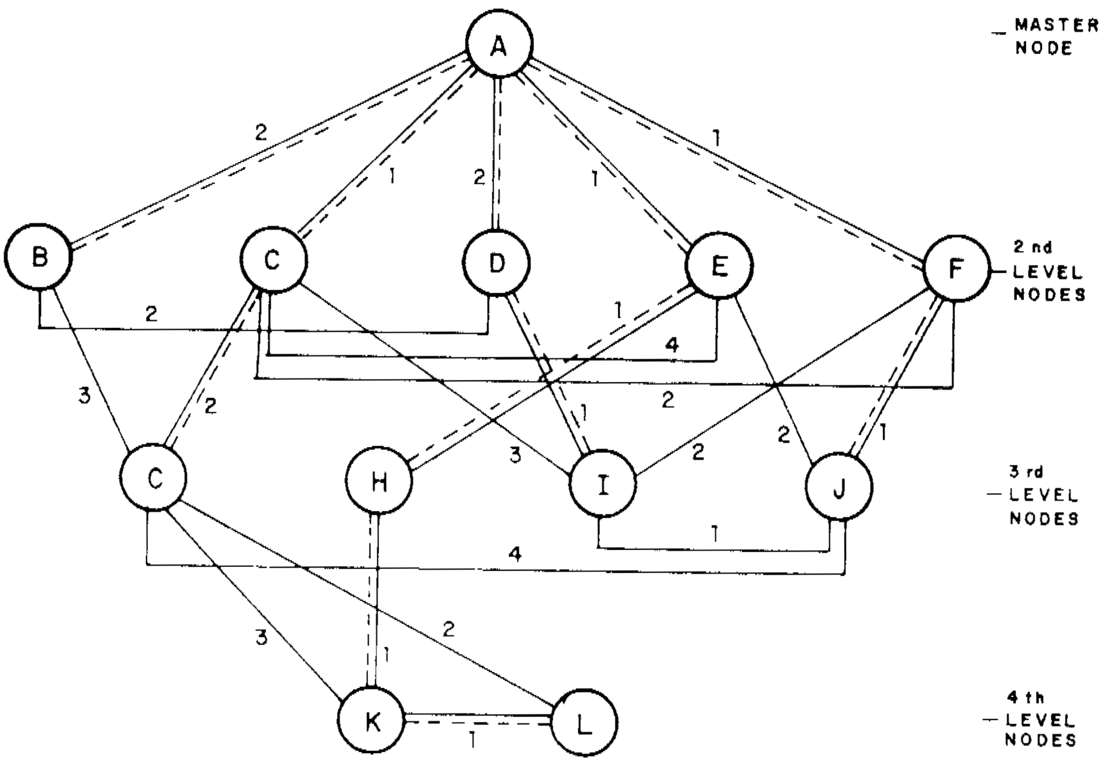

NUMBERS BY LINK LINES ARE INACCURACIES (VARIANCES)

---- DASHED LINES REPRESENT BEST TIME REFERENCE PATHS

FIGURE 2. EXAMPLE NETWORK OF FIGURE 1 ARRANGED IN A HIERARCHY WITH NODE A AS MASTER

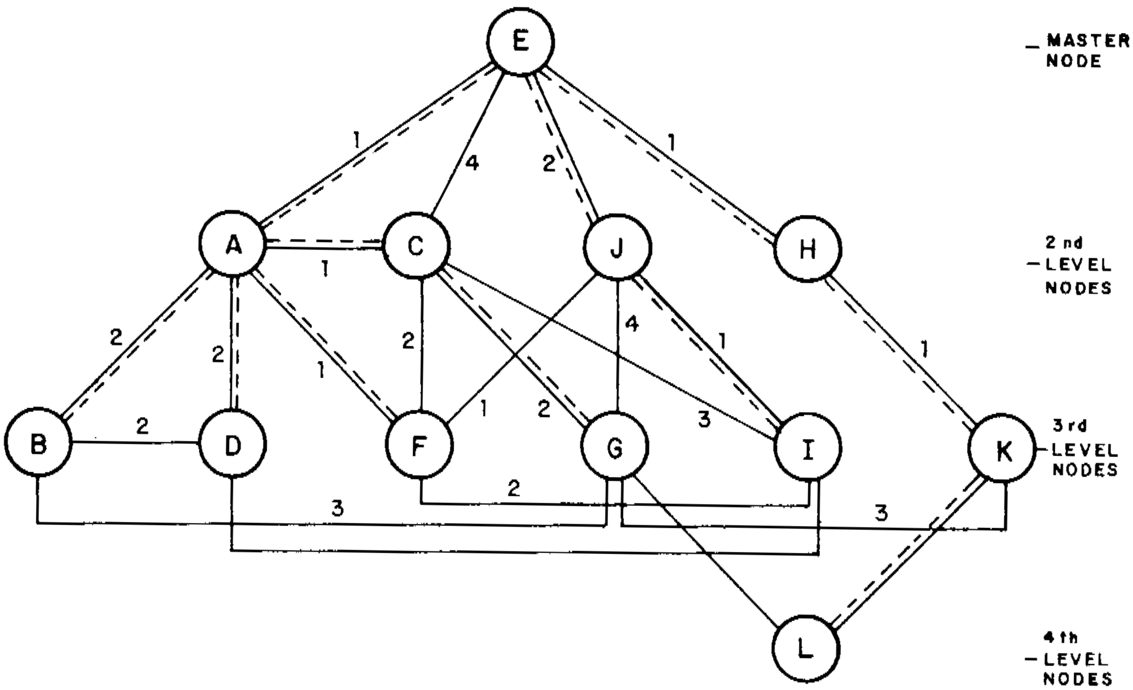

NUMBERS BY LINK LINES ARE INACCURACIES (VARIANCES)

----- DASHED LINES REPRESENT BEST TIME REFERENCE PATHS

FIGURE 3. EXAMPLE NETWORK OF FIGURE 1 ARRANGED IN A HIERARCHY WITH NODE E AS MASTER

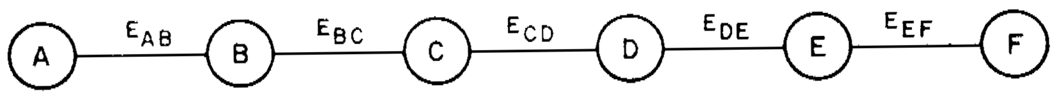

FIGURE 4. TANDEM LINK EXAMPLE

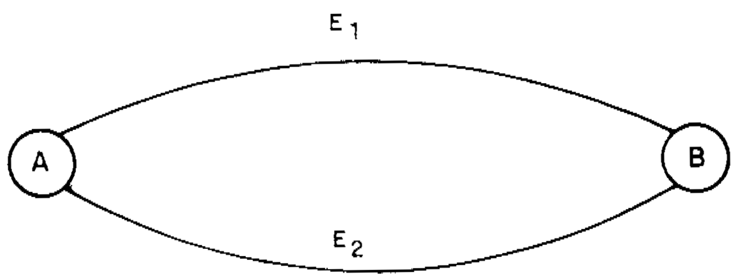

FIGURE 5. PARALLEL LINK EXAMPLE

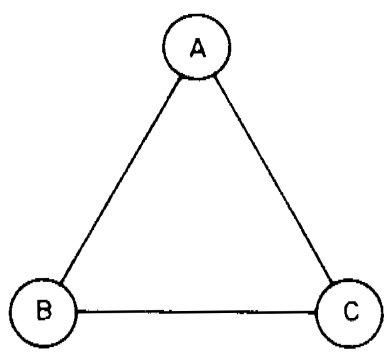

FIGURE 6. EXAMPLE THREE NODE NETWORK

ACKNOWLEDGEMENT

The author wants to thank Dr. Eugene P. Kaiser, Mr. Allen C. Trainer, and Mr. Franklin R. Wiesman for their editorial suggestions, and Mrs. Anita J. Hughes for preparing the manuscript.

REFERENCES

1. DCASEF TC 39-73, "Time Reference Concept for the Timing and Synchronization of the Digital DCS", H. A. Stover, July 1973.

2. "A Time Reference Distribution Concept for a Time Division Communications Network", Proceedings of the Fifth Annual NASA Department of Defense Precise Time and Time Interval (PTTI) Planning Meeting, NASA, Goddard Space Flight Center, Greenbelt, Maryland (December 4-6, 1973), H. A. Stover, December 1973.

3. "Coordinated Universal Time (UTC) As A Timing Basis for A Digital Communications Network", H. A. Stover, EASCON 1974 Record.

4. DCEC TR 43-75, AD-A021934, "Communications Network Timing", H. A. Stover.

5. "Synchronization in a System of Interconnected Units," Darwin and Prim, U.S. Patent No. 2,986,723, May 30, 1961.