Methodologies For Steering Clocks

H. Chadsey, USNO

Abstract

One of the concerns of the PTTI community is the coordination of one time scale with another. This is accomplished through steering one clock system to another, with a goal of a zero or constant offset in time and frequency. In order to attain this goal, rate differences are calculated and allowed for by the steering algorithm. This paper will present several of these different methods of determining rate differences. Ideally, any change in rate should not cause the offset to change sign (overshoot) by any amount, but certainly not by as much as its previous absolute value. The advantages and disadvantages of each depend on the user's situation.

INTRODUCTION

Although control system theory is not new, it has been more highly developed in some fields than others. One good example is rocket science and the degree of control theory development needed for the US Space Shuttle to "catch" the Hubble Space Telescope in December of 1993. The trick was to meet up with the orbiting unit and grab it. If the control was wrong, the Shuttle would not reach the unit; pass by it; or, worse, crash into it. This approach to an offset (in position) is done through the precise firing of rockets. It is a very critical operation because the rockets have only two states: off and full thrust.

The same type of problem faces a laboratory trying to steer the frequency of a clock. The objective is to maintain a clock at zero (or some other fixed) time offset from some reference clock. While steering to the desired value, the offset should not be allowed overshoot by any amount, but most certainly not by the same or more than its previous maximum offset. An important difference is that more precise clock control can be obtained because a variable steer rate algorithm can be determined. This can be done by taking the principles of control theory as applied in other fields of operation and applying them to the control of clocks.

REASON TO STEER

The need for steering may be understood by looking at what happens when no steering is attempted. A lab might only monitor the time offset of the clock to some standard periodically (hourly, daily, weekly, etc.). From these periodic measurements, missing values would be derived by interpolation, operations would be carried out with no controlling of the clock. For some applications, especially if they are performed over short periods of time, this is acceptable. However, all clocks have a rate which changes. This is called drift and it is not constant. As a result, the clock will be very far ahead or behind and vary in offset amount when compared to the reference clock, creating problems for some operations. (See Figure 1.)

STEERING METHODS

To correct for the drift of a clock, the most rudimentary of steering methods can be used. For example, the clock may be closely monitored and allowed to increase the offset value as compared to the standard. Once the clock has reached a predetermined offset, it may be time–stepped to a different value. Operations continue by using the varying offset values and interpolating between them as needed. This is much like the example of the Space Shuttle cited earlier, where the space craft is allowed to drift, controlled through the use of varying–length, full–throttle corrective actions. In the controlling of cesium clocks with this method, there are two major potential problems. First, operations can be disrupted when the clock is time–stepped. This can in some cases be avoided by performing the steps at times when the clock is not being used for operations. When the clock is adjusted, very close monitoring must be performed and methods developed to determine values of offset during the stepping procedure. Second, cesiums and many other types of clocks can have their characteristics changed when they are time–stepped. Cesium clocks have been known to change their drift rate when adjustments of any type are made to them. Again, it depends on how the clocks are being used whether this will have an adverse effect on operations. (See Figure 2.)

We turn now from the manipulation of the clock to controlling of the output from the clock. Timing is controlled, not by adjusting the clock itself, but through adjustment of its output with a phase microstepper or similar device.

The most efficient and drastic of these steering adjustments is commonly referred to as the "Bang–Bang" mode of operation. The crudest form is the two–stage steering algorithm. This is the method currently employed by the GPS Master Control Station to control GPS time. This methodology lets the clock(s) drift at its natural rate until a predetermined offset is reached. At that time, a frequency change is made to the output (using a phase microstepper, or adjustment of the clock, etc.). The new drift rate of the output is in the opposite direction and at a greater rate than the natural drift of the clock. This new rate is kept until the clock reaches another predetermined offset value, when the rate is again changed back to its first value. These rates are currently $\pm 1.0 \times 1-^{-19}$ seconds per second squared for the GPS system. As a result of this two value steering, the clock oscillates between the extreme offset values. Because the natural rate of the clock is in only one direction, the "wave" is asymmetric. (See Figure 3.)

The next step in complexity is the three–stage steering algorithm. Here, one has the additional state of a zero rate of steer to the system. The theory is that while the clock offset is within a narrow range of offset, a zero rate of steer is employed by the system. When the offset is outside this narrow range, the upper or lower steer rate is employed to move the clock to within the narrow offset range when the zero rate is again used. While this is very easy to perform programmatically, it still does not correct for the natural rate of the clock. It can also produce a wave pattern which may not be stable enough for the operational use of the clock output. (See Figure 4.)

The previous steering methods use a fixed, predetermined rate of adjustment. We next consider the possibilities of enhancement to a system when a variable rate adjustment is implemented. First, when the rate is not fixed, we will discuss how it is determined. Second, we will see how the rate correction is applied.

One of the first conclusions is that the clock can now be adjusted for its natural drift. We no longer have to hassle with an asymmetric wave of offsets. But how can the natural drift rate be determined?

One method would be to subtract the first data point collected after the last rate change was applied from the last data point collected and divide by the number of days in the interval between them. This would result in a rate per day which can easily be used to calculate an adjustment. This method, however, does have a potentially large fault. If the data are noisy, the rate determined could be of the wrong magnitude and/or sign.

A second method would be to take an average of the differences between successive days of data. This would reduce the likelihood of problems. For a well–behaved system, the taking of successive differences will allow one to construct where the next data point will lie. As a result, a rate change can be determined and tested before it is applied to the system. The problems here are that, depending on the precision of measurement and control desired, the clock system may not be a well–behaved system; large amounts of data are required which are not always available; and this method does not react well to sudden change, such as a clock jump.

A third method would be to perform some type of data analysis on the data points in order to find the slope of the values. A system of tracking the rate of change of a moving average or linear line fitting with slope determination over short time periods does very well. It is improved when, on larger data sets, a data filter is used to remove outlying points. After some testing, I chose the linear fit method because of its more direct approach, its flexibility to filter outlying points, and the fact that "one cannot design a filter better than the optimum linear filter".

Now that we have a way of determining the natural drift of the clock, we can use that to help in controlling it. This leads to a methodology I will refer to as graduated steering. It is graduated because the rate of correction applied is no longer fixed, but varies according to some algorithm. Because we can determine the natural drift of the clock by one of the methods previously discussed, it follows that we should use that as a starting point for the amount of the change.

For example, if the rate of the clock was found to be positive 30 nanoseconds per day, then we can apply a rate of -30 nanoseconds per day by use of a phase microstepper. The result would be a clock with a zero rate of change as compared to the reference standard. (The change can also be an adjustment of the C–field, but a phase microstepper allows for finer adjustments.)

Now that we have created a clock with a zero drift, we must get a zero offset to the reference clock in order for our operational requirements to be met. This could be done by time stepping ("banging") it back to zero. This would require the calculation and manual operation of the time step. It would also mean that our operation could be disrupted by the time step. The long-term complication would be that if the natural rate of change for the clock was not determined exactly, after a period of time the procedure would have to be repeated. This requires constant monitoring of the system, training of personnel in how to make the corrections, and ascertaining that the corrections are made when they are needed.

The correction to a zero (or near zero) offset can be performed by modifying the rate determination program to also add a small amount to that required to achieve zero drift. But, how should this additional amount be determined? A method that was used at USNO based the additional amount on the clock drift. This was found to lead to some undesired effects. Using the offset value as a factor for additional calculation provides much better control of the clock. The idea of steering back to a zero offset should be based not only on the rate at which one is moving toward or away from it, but also on the present offset value. If, for instance, one's clock has a rate of 50 nanoseconds per day and the current offset is zero, one would then introduce a rate change of -50 nanoseconds per day to achieve zero offset. On the other hand, if one's clock has a rate of 50 nanoseconds per day and the present offset value is 100 nanoseconds, one would introduce a rate change of -50 nanoseconds per day to correct for the rate of the clock plus an additional amount to get back to zero offset. The additional amount could be the present offset value divided by some damping factor, say 4. In this case, the offset would be zero in 4 days time and another rate correction introduced to flatten the rate of the clock. This method of simple graduated steering will correct for the natural drift of the clock and if the clock is offset from a standard reference, it will adjust it back to zero offset. The USNO Master Clock is steered in a similar fashion.

Under ideal circumstances, this simple graduated steering method works fairly well. The complication arises when one is unable to make one or more of the rate corrections. This can be caused by the failure of one computer to instruct another controlling the microstepper to make a change (e.g. because of a faulty phone connection), by computer or microstepper failure, or, if the operation is performed manually, by the absence of required personnel. The magnitude of this error can be lessened if a damping factor is used or if the one in use is increased.

PRESENT USNO STEERING METHOD

The present USNO method of steering remote clocks has developed into a very sophisticated process that is totally computer-controlled. Operator intervention is needed only in case of equipment failures or other extraordinary events (e.g. clock jumps or clock replacement). The process begins by a program determining the offset of the remote clock from the USNO Master Clock using GPS. This can be done using one of several methods. USNO currently uses the 48-hour running linear-fit melting-pot method. Once the offset for a series of days has been determined, the steering rate determination program begins to work.

The program to determine an adjustment to the phase microstepper performs several tests before the calculation is actually performed. The first test is to determine whether or not a steer correction is permissible. There is a big danger in calculating rate changes using data that contains a mix of data from before and after implementation of the last rate correction. Calculation using mixed data can lead to the steering process causing the clock offset values to oscillate in a very extreme manner. (See Figure 5.) Because of this danger and the use at USNO of the two-day fit method of GPS data processing, the steer rate determination does not use any offset values two days following the implementation of the last correction.

The next test is to make sure that there are enough data to determine an accurate rate of the remote clock. This requires at least two days of data. The more data collected, the more accurate the determination will be. However, if too many days are used, problems such as fast reaction to clock jumps and clock replacement will be created. So, when possible, the minimum amount of days between steers is used.

As a result of the these first two tests performed by the program, rate changes to the remotely controlled clocks occur no sooner than once every four days. Combination of the avoidance of mixed data and use of sufficient data points can improve clock stability by a factor of ten.

The actual rate change is then determined. From the above tests and predictions of the offset, it is easy to determine whether the clock is moving toward or away from zero offset and whether the offset changed sign during the time between the last rate correction and the present or will change sign before the time of the next rate correction. Programmatically, this creates a four-state test switch:

1. All collected and predicted values are of the same sign and the trend is toward a zero offset. In this case, the steer rate change would be the predicted value divided by the damping factor and again divided by the number of days between steer rate changes. This double division prevents overcorrection of the clock.

2. All collected and predicted values are of the same sign and the trend is away from a zero offset. In this case, the rate change would be the rate of the clock away from zero to flatten the clock rate plus an additional amount to direct the clock back toward zero. This additional amount would be the predicted value divided by the damping factor. The number of days between steers is NOT used in the divisor, as this would not cause the clock to turn around its direction of offset travel. Overcorrection is kept to a minimum because of the time intervals at which events occur. The correction is applied when the clock has already moved further away from zero than the collected data indicate.

3. All collected values are of one sign and the predicted value is of the opposite sign. In this case, it was found from the many possible ways it can occur that the best correction is to make the rate change equivalent to the rate of clock, thus zeroing the clock rate.

4. Some of the collected values are of one sign and the rest of the collected values and the predicted value of the opposite sign. In this case, the rate change would again be the rate of the clock away from zero to flatten the clock rate, plus an additional amount to direct the clock back toward zero. This additional amount would be the predicted value divided by the damping factor. The number of days between steers is NOT used in the divisor, as this would not cause the clock to turn around its direction of offset travel. Overcorrection is kept to a minimum because of the time intervals at which events occur. The correction is applied when the clock has already moved further away from zero than the collected data indicate.

The correction is then placed in units of nanoseconds rate of change per day. Another program implements the change in the phase microstepper at the remote unit. The process is complete until the next correction is needed.

RESULTS OF STEERING

Figures 6 and 7 show the results of the implementation of the present control program for two sites. The clock being controlled in Figure 6 is located in the office area of an Air Force PMEL. The clock being controlled in Figure 7 is located in a controlled environment chamber. Both are using the same program for determining the rate adjustment for the phase microstepper.

CONCLUSION

With some thought and investigation into the control theories used in other fields of operation, the timing community can develop programs to provide more accurate and precise control of time references. This can only lead to improvements in operations and an ultimate savings of money and personnel time. Control of a clock or system of clocks can range from a simple manual operation concerning the periodic readjustment of a clock to an elaborate computer program control operation. The computer control program can be as complicated as the programmer desires. "...intuition is at a premium in nonlinear design" and the control of a clock system in a changing environment is definitely a nonlinear operation. Beyond a certain point of complexity, the more that is added to a program, the more likely a control error will occur.

USNO minus UNSTEERED CLOCK

(Figure 1)

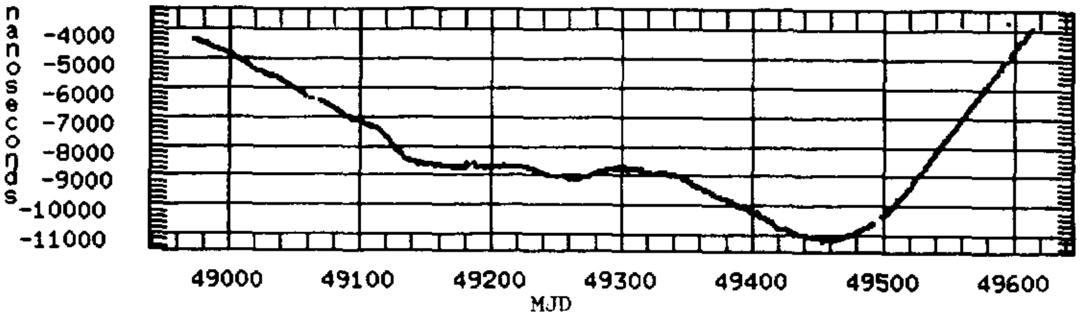

USNO minus STEPPED CLOCK

(Figure 2)

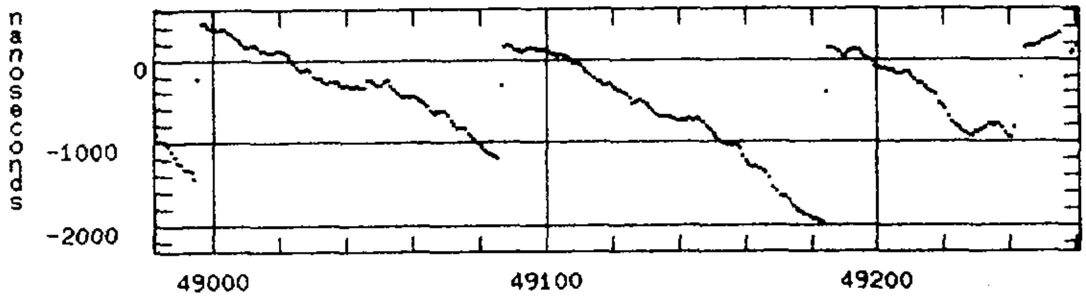

REFERENCE minus CLOCK USING TWO-STAGE STEERING

(Figure 3)

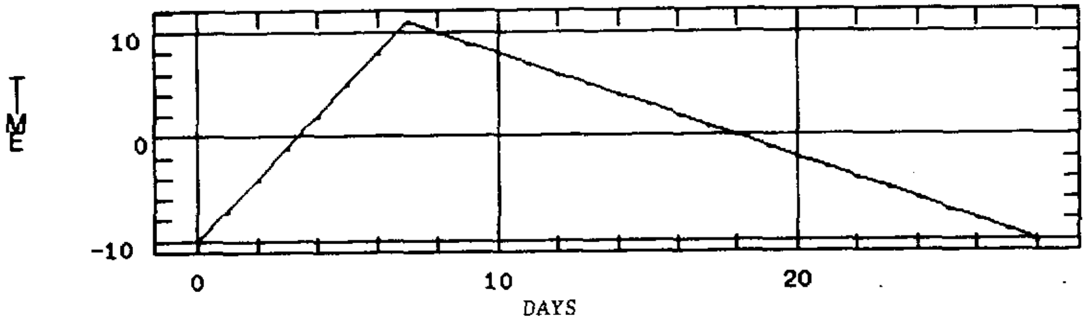

REFERENCE minus CLOCK USING THREE-STAGE STEERING

(Figure 4)

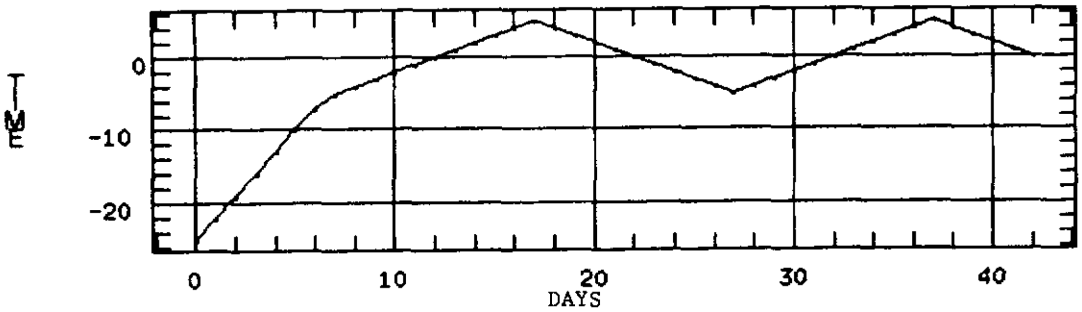

USNO minus MIXED DATA STEERED CLOCK

(Figure 5)

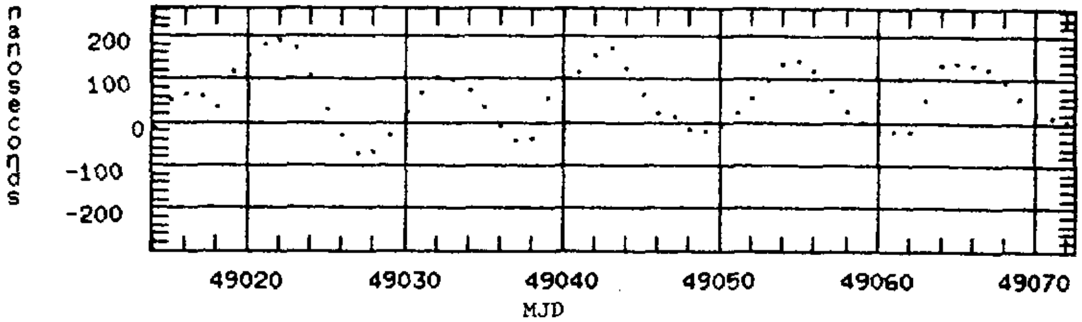

USNO minus STEERED CLOCK IN OFFICE

(Figure 6)

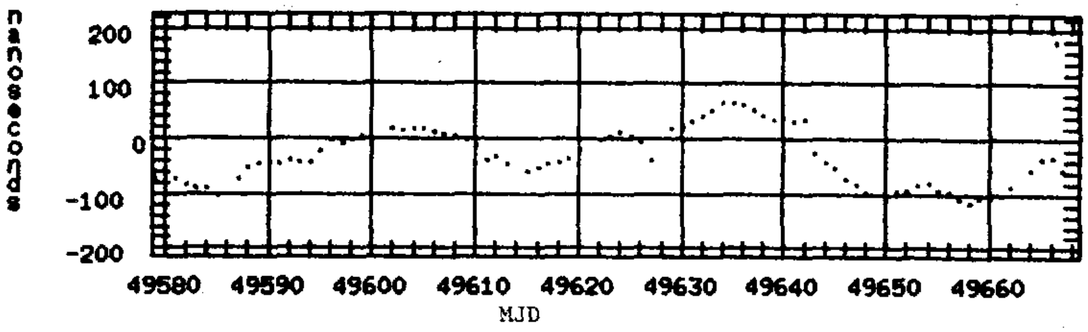

USNO minus STEERED CLOCK IN CONTROLLED AREA

(Figure 7)