Time Scale Algorithms for an Inhomogeneous Group of Atomic Clocks

C. Jacques, J.-S. Boulanger, R.J. Douglas, D. Morris,
S. Cundy, and H.F. Lam
Time and Frequency Standards Group
Institute for National Measurement Standards
National Research Council of Canada
Ottawa, CANADA K1A 0R6

Abstract

Through the past 17 years, the time scale requirements at NRC have been met by the unsteered output of its primary laboratory cesium clocks, supplemented by hydrogen masers when short-term stability better than $2 \times 10^{-12}\tau^{-1/2}$ has been required. NRC now operates three primary laboratory cesium clocks, three hydrogen masers and two commercial cesium clocks.

NRC has been using ensemble averages for internal purposes for the past several years, and has a real-time algorithm operating on the outputs of its high-resolution (2×10^{-13} s @ 1 s) phase comparators. The slow frequency drift of the hydrogen masers has presented difficulties in incorporating their short-term stability into the ensemble average, while retaining the long-term stability of the laboratory cesium frequency standards.

We report on this work on algorithms for an inhomogeneous ensemble of atomic clocks, and on our initial work on time scale algorithms that could incorporate frequency calibrations at NRC from the next generation of Zacharias fountain cesium frequency standards having frequency accuracies that might surpass 10^{-15}, or from single-trapped-ion frequency standards ($Ba^{+}, Sr^{+}, \ldots$) with even higher potential accuracies.

We present and discuss the requirements for redundancy in all the elements (including the algorithms) of an inhomogeneous ensemble that would give a robust real-time output of the algorithms.

Introduction

Time scale algorithms are of wide interest because an ensemble average should behave better than the best clock of the ensemble, in terms of accuracy, stability and reliability: to some extent the misbehavior of any clock of the ensemble can be filtered out without disturbing appreciably the ensemble average. The stability of a new clock can be added to the ensemble average without major disturbance.

Historically, at NRC, the implementation of a time scale algorithm at NRC has been delayed by the excellent behavior of the laboratory cesium clocks of the Cs V and Cs VI designs. Used alone, they met most of the needs of TA(NRC) and UTC(NRC), with the other needs being met by NRC's masers. The advantages of developing a good time scale algorithm were not as evident as the advantages of maintaining a "good clock".

Nonetheless, a time scale algorithm was developed for the NRC ensemble, and used to assist in primary clock evaluations. It was particularly beneficial for evaluations when the NRC ensemble was operating in a degraded mode caused by problems either with our building's environment or with the clocks themselves.

This algorithm was developed to exploit the short-term characteristics of the NRC clock ensemble for the purposes of clock diagnostics, initially to run with the NRC three-channel phase comparators at 10^{-12} s resolution, taking phase data every 10 s. With the new NRC masers, a higher resolution phase comparator was desirable, and the old design [1] was improved to give a unit which gives phase differences each second (averaged over one second) with 0.1 ps resolution, and a short-term noise level of less than 0.2 ps. We have built, and are commissioning two phase comparator systems for full redundancy. Each phase comparator system can record 32 phase differences taken from up to 16 5 MHz sources, and has standard serial output at 9600 bps. The main data acquisition software runs on MS-DOS computers under Deskview. We wrote a driver to accept and check the phase comparator data, and write the files to disk. The data acquisition computer system also runs the ensemble average software in real-time (with less than 1.5 s delay from any phase step to the ensemble average: up to 1 s delay in the phase comparator, up to 0.3 s in serial transmission, and up to 0.2 s for the recalculation of the ensemble average). The same computer system flexibly displays the phase difference results in a wide variety of formats.

Other operational advantages of real-time time scale algorithms also became evident at NRC as automated inter-laboratory time transfer with resolutions of better than 200 ps came on line in our laboratory. Also a real-time algorithm promises advantages during evaluations or during repairs to the primary laboratory cesium standards.

An algorithm seems attractive for exploiting the superior timekeeping characteristics of NRC's two new masers, and could gracefully incorporate frequency calibrations from the initially short periods of operation of the new cesium frequency standards — cesium Zacharias fountains — and single-ion frequency standards, which are under development at NRC and will likely have frequency accuracy capabilities in the 10^{-15} range. An algorithm based on Kalman filtering seems to be a promising candidate to meet our requirements of good stability in both the short term and long term while keeping the accuracy that cesium beam frequency standards and their successors provide.

Kalman Filtering

Of the many time scale algorithms used in different institutions and laboratories [2, 3, 4, 5, 6, 7], an algorithm based on Kalman filtering seemed the most promising to meet our needs. It is an optimum estimator in the minimum squared error sense, it applies to dynamic systems with a proper adaptive filtering technique, it has the optimum transient response, and it can address the requirements of both short and long term stability.

We have used Kalman filtering at NRC for several years, for algorithm evaluation purposes although the ensemble average has been a very helpful tool for clock evaluations from early in the development. We have set up a full project to pursue the development of the algorithms, and fine tune an algorithm that we envisage will include rejection and correction of outliers for both phase and frequency jumps.

The purpose of this paper is not to develop the Kalman filter equations. This development can be found in many good books[8, 9]. Briefly, if we have a linear system with the state-space description

$$\begin{cases} \mathbf{x}_{k+1} &= A_k \mathbf{x}_k + \xi_k \\ \mathbf{v}_k &= C_k \mathbf{x}_k + \eta_k \end{cases} \tag{1}$$

where A_k, C_k are $n \times n$, and $q \times n$ constant matrices, respectively the state transition matrix and the connection matrix between the measurement $\mathbf{v}_k$ and the state vector $\mathbf{x}_k$, and $\{\xi_k\}$ and $\{\eta_k\}$ are respectively system and observation noise sequences, with known statistical information such as mean, variance, and covariance, we can derive the following Kalman filtering process:

$$\begin{cases} P_{0,0} &= \mathrm{Var}(\mathbf{x}_0) \\ P_{k,k-1} &= A_{k-1} P_{k-1,k-1} A_{k-1}^\top + Q_{k-1} \\ G_k &= P_{k,k-1} C_k^\top (C_k P_{k,k-1} C_k^\top + R_k)^{-1} \\ P_{k,k} &= (I - G_k C_k) P_{k,k-1} \\ \hat{\mathbf{x}}_{0|0} &= \mathrm{E}(\mathbf{x}_0) \\ \hat{\mathbf{x}}_{k|k-1} &= A_{k-1} \hat{\mathbf{x}}_{k-1|k-1} \\ \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + G_k (\mathbf{v}_k - C_k \hat{\mathbf{x}}_{k|k-1}) \\ k &= 1, 2, \ldots \end{cases} \tag{2}$$

where

$P_{k,k-1} = \mathrm{Var}(\mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1})$ the variance of the difference between the state vector and the prediction $\hat{\mathbf{x}}_{k|k-1}$,

$P_{k,k} = \mathrm{Var}(\mathbf{x}_k - \hat{\mathbf{x}}_{k|k})$ the variance of the difference between the state vector and the estimation $\hat{\mathbf{x}}_{k|k}$,

$Q_k = \mathrm{Var}(\xi_k)$,

$R_k = \mathrm{Var}(\eta_k)$.

Equation 2 is a recursive scheme that, when applied to the incoming data $\mathbf{v}_k$, produces predictions $\hat{\mathbf{x}}_{k|k-1}$ and estimations $\hat{\mathbf{x}}_{k|k}$ of the state vector $\mathbf{x}_k$. The difference $\mathbf{v}_k - C_k \hat{\mathbf{x}}_{k|k-1}$ is called the innovation. It can be shown that

$$\mathrm{Var}(\mathbf{v}_k - C_k \hat{\mathbf{x}}_{k|k-1}) = C_k P_{k,k-1} C_k^\top + R_k. \tag{3}$$

For a description of a clock, we need to know its phase and frequency; in some case the frequency aging is needed. Since absolute time is not known, the exercise is equivalent to comparing two or more clocks together: that means that the state vector components are the phase difference $x(t)$ and the relative frequency $y(t)$. The following equations are for the phase difference and relative frequency between two clocks. A generalization to more clocks is straightforward.

$$\mathbf{x}(t) = \begin{pmatrix} x(t) \\ y(t) \end{pmatrix} \tag{4}$$

and

$$\hat{\mathbf{x}}_{k|k-1} = \hat{\mathbf{x}}(t \mid t - \delta) \tag{5}$$

where $\hat{\mathbf{x}}(t \mid t - \delta)$ is the prediction on $\mathbf{x}(t)$ given $\hat{\mathbf{x}}(t)$ from $t = 0$ up to $t - \delta$. The state transition matrix is:

$$A_k = \begin{pmatrix} 1 & \delta \\ 0 & 1 \end{pmatrix}; \tag{6}$$

and since we measure the phase difference between two clocks, the connection matrix between the measurement and the state vector is

$$C_k = \begin{pmatrix} 1 & 0 \end{pmatrix}. \tag{7}$$

R_k is the variance of the measurement error:

$$R_k = \sigma^2_{\mathbf{v}_x}. \tag{8}$$

Now, if we have a careful look to the equations 2 of the Kalman filtering process, we see that the second, third and fourth equations are self-sufficient as a group: the matrices $P_{k,k-1}$, $P_{k,k}$ and G_k are updated independently of the data $\mathbf{v}_k$! There is nothing wrong with that if the description of the model is right, and if the behaviour of any clock in the ensemble is unchanging. The latter is practically impossible to realize, at least on a long term basis. The development of adaptive filtering is needed in our case. This means that the parameters of the model, specifically the variance matrix Q_k of the system noise, must be continuously evaluated. The frequency of evaluation depends on the type of noise process involved.

Although we are still evaluating different adaptive filtering approaches, the one developed by Stein[5] looks the most appropriate of all the adaptive filtering approaches we have studied until now[9, 10, 11]. The approach is based on the fact that the variance of the innovation, equation 3, links the measurement $\mathbf{v}_k$ with $P_{k,k-1}$, which is defined in terms of Q_{k-1} in equation 2. Our next step will be to implement a slightly modified version of Stein's approach, the modifications are more on the procedure of calculating the parameters rather than on the underlying principles.

Results using the current algorithm

We present here an analysis of our internal time scale algorithm based on Kalman filtering by looking at two different sets of clock data — relative phase measurements between six pairs of clocks taken at 10 minute intervals — taken respectively between MJD=48795 and 48865 for 70 days (from 92-06-22 to 92-08-31), and between MJD=48895 and 48915 for 20 days (from 92-09-30 to 92-10-20). These two sets of clock data correspond to periods of relatively stationary behaviour of the clocks. Of the different clocks included in the calculation of the time scale, there are the two new NRC hydrogen masers, H3 (H) and H4 (h), on which there are still experiments done, often resulting in frequency steps; for that reason, H3 has been included in the calculation of the first period of 70 days for only 37 days, between MJD=48795 and 48832. H4 is not included at all in the 70 day period. Besides the two hydrogen masers, the time scale algorithm is calculated from the relative phase measurement of three primary laboratory cesium beam atomic clocks, Cs V, Cs VI A and Cs VI C, and two commercial cesium beam atomic clocks P and p (HP5061-A's with super tubes). The six phase differences were taken with two three-channel phase comparators with 1 ps resolution, and the calculations were done using phase difference data taken every 10 minutes.

The time scale algorithm evaluation includes also Cs VI C (undergoing a full evaluation during part of this period) and the two commercial cesium clocks. For that period, though, their measured stabilities were not good enough to give them important weights in the calculation, and we will not present results for them.

To analyze the time scale algorithm, we present the Allan deviation graphs of each of the other clocks vs the ensemble. As expected from Figs. 1-3, the hydrogen masers H show a very good short term stability that is reflected in the algorithm. Fig. 1 shows the Allan deviation of free- running maser H3 vs the ensemble. In Figs. 2 and 3 the Allan deviations are for the masers under cavity servo control. In Figs. 1-5, the upper and lower traces show the limits of the interval of confidence (95 %) in the evaluation of the Allan deviation. After a week, Cs V starts to take over when it reaches the $\sigma_y(\tau) = 2 \times 10^{-14}$ level, Fig. 4. A longer term analysis, of the order of 1 year, would show a larger contribution of Cs V, and also of Cs VI A and Cs VI C, since their long term stability is better than the hydrogen masers. A comparison of Figures 1, 4 and 5 shows that the cesium clocks and the hydrogen maser H3 are at the same level of stability after a week, defining the beginning of a cross-over region before the long term stability of the hydrogen maser deteriorates. At this level of stability, $\sigma_y(\tau) = 2 \times 10^{-14}$, and on a period of 70 days, it is easy to appreciate the difficulties of determining from the Allan deviation over this period which clock(s) should be pulling the algorithm the most. This time scale algorithm allows a maximum weight of 0.8 for the contribution of the best clock.

Other long-term questions have not been resolved in this algorithm. It does have built-in consistency from the continuous evaluation of the weights from predictions and estimations for the calculation of the phase of the ensemble. However, the same standard of consistency is not implemented for frequency, nor for aging or related effects. The consequence is that the long term behavior of the time scale is not optimally controlled, with phase comparisons taken every ten minutes as clock data. The algorithm could be pulled by the hydrogen masers even if the long term stability of a cesium clock is better. This problem will be addressed in the further developments in our time scale algorithm.

Additional requirements for a real-time system

The results of our experience with ensemble averaging, such as that presented above, has encouraged us to undertake a project aimed at implementing a real-time algorithm for UTC(NRC). Our old system has used the proper time output of our "best" primary cesium clock, as PT(NRC), adjusted for our 100 m elevation by a microstepper to convert to UTC(NRC). We have examined the requirements of a time scale system, and we are building a system shown in Fig. 6. It measures phase differences each second and controls a quartz oscillator giving UTC(NRC) by calculating and outputting a correction to the quartz oscillator frequency with a delay of less than 1.5 s. Most components of the system have on-line backup, both to minimize the effects of component failure and to simplify component maintenance and upgrading. Most of the hardware has been operating for over a year, and the algorithm presented above has been re-implemented on the redundant PC's (switching from the Hewlett-Packard Basic (Rocky Mountain Basic) and Infotek compiler to Microsoft's Quickbasic and Basic System 7 for MS-DOS 80486 computers). We have encountered more difficulties than expected with compiler errors as well as the normal coding errors. We are in the process of independently re-coding the algorithm in FORTRAN, both for speed and as a further check on the coding and compilers: we plan to use the greater portability of the FORTRAN code to run the algorithm on VAX VMS, IBM VM, Silicon Graphics UNIX as well as PC's under MS-DOS. The greater speed will expedite the offline examination of different strategies for algorithms running on our small, inhomogeneous ensemble of atomic clocks.

When one examines the behaviour of clocks, what looks simple is in fact more complicated. The stability of a clock, mostly evaluated by the Allan variance, varies with the time interval on which it is evaluated. If a phase or frequency step happens, it must be not only detected but also determined as being either phase or frequency. A time scale algorithm must deal with changes in stability, both in time and in time interval, and in phase and frequency. The evaluation of the parameters of the algorithm, the system noise and the weights of the clocks, must be as automatic as possible without loss of control. A well designed algorithm lets the human operator keep track of all the evaluated parameters to keep control in case the algorithm behaves inappropriately. The reliability is one of the aspect on which the algorithm will be evaluated extensively. Initially we expect to be running different real-time algorithms in parallel, constrained to lie within a time window around the old "best clock" algorithm. The time window would have to be manually set and reset when necessary.

Conclusion

Until now, UTC(NRC) has been derived from one of the three laboratory primary cesium beam clocks, Cs V, Cs VI A and Cs VI C. Their excellent capabilities has delayed the implementation of a time scale algorithm. The advantages of a time scale algorithm for NRC's inhomogeneous group of atomic clocks were not as evident as the advantages of maintaining a "good clock".

We presented in this paper some ideas for the development of a timescale algorithm based on Kalman filtering. The choice of Kalman filtering is dictated by our requirement for both short and long term stability of the ensemble. By taking phase differences between pairs of clocks instead of time of clocks, we make sure that the different variance matrices involved in the recursive calculation of the time scale do not diverge. The system noise variance matrix must also be updated dynamically to take into account the change in behavior of any clock of the ensemble. Continuity, not only in phase but also in frequency, and aging if measured, must be implemented in the algorithm. This is to avoid pulling of the algorithm by clocks which have the best short term stability, like the hydrogen masers, but poorer long term stability. We have presented an analysis of the time scale algorithm, used internally only, for two periods of time of 20 and 70 days respectively. The algorithm was found to give the most weight to the best clocks in the short term, but the analysis didn't allow us to evaluate the algorithm on the long term.

We considered also the requirements of an on-line time scale system generating UTC(NRC), from the measurement of phase differences to the control of the quartz oscillator. Many of the components will be redundant, the measurement system and the computer calculating the algorithm, to detect any phase comparator or computer error. The program and its coding will be checked against different source of errors like coding, compiler, and the algorithm itself. For that matter, the coding will be checked on different computer architectures, not only on a PC DOS 80486 where it will be implemented.

In the future, a time scale algorithm will have be optimized to facilitate exploiting the initially short periods of operation of NRC's new frequency standards: cesium Zacharias fountains and single-ion frequency standards, which will likely have frequency accuracy capabilities in the 10^{-15} range. Unless there are surprising advances in time intercomparisons, the promises of good algorithms will have to be realized to allow inter-laboratory frequency intercomparisons at this level. Subsequent work will focus on the optimal inclusion of the intervals of operation of these higher accuracy frequency standards into the ensemble average of our inhomogeneous group of atomic clocks.

References

[1] D.W. Allan, and H. Daams, Proceedings of the 29th Annual Symposium on Frequency Control, Fort Monmouth, New Jersey, U.S.A., 1975, pp. 404-11

[2] P.A. Clements, and B.P. Gibbs, in *Third International Time Scale Algorithm Symposium*, Turin, Italy, 1988, pp. 289-310

[3] M. Mnackri, and C. Thomas, in *Third International Time Scale Algorithm Symposium*, Turin, Italy, 1988, pp. 313-26

[4] S.R. Stein, in *Third International Time Scale Algorithm Symposium*, Turin, Italy, 1988, pp. 345-57

[5] S.R. Stein, Research and Development Technical Report SLCET-TR- 89-0717-F, US Army Laboratory Command, 1989

[6] P. Tavella, and C. Thomas, Metrologia **28**, 57-63 (1991)

[7] M. Weiss, and T. Weissert, Metrologia **28**, 65-74 (1991)

[8] C.K. Chui, and G. Chen, *Kalman Filtering with Real-Time Applications*, Second Edition, Springer-Verlag, 1991

[9] R.G. Brown, *Random Signal Analysis and Kalman Filtering*, John Wiley & Sons, 1989

[10] D.T. Magill, IEEE Transactions on Automatic Control AC-10(4), 434-9 (1965)

[11] R.K. Mehra, IEEE Transactions on Automatic Control AC-15(2), 175-84 (1970)

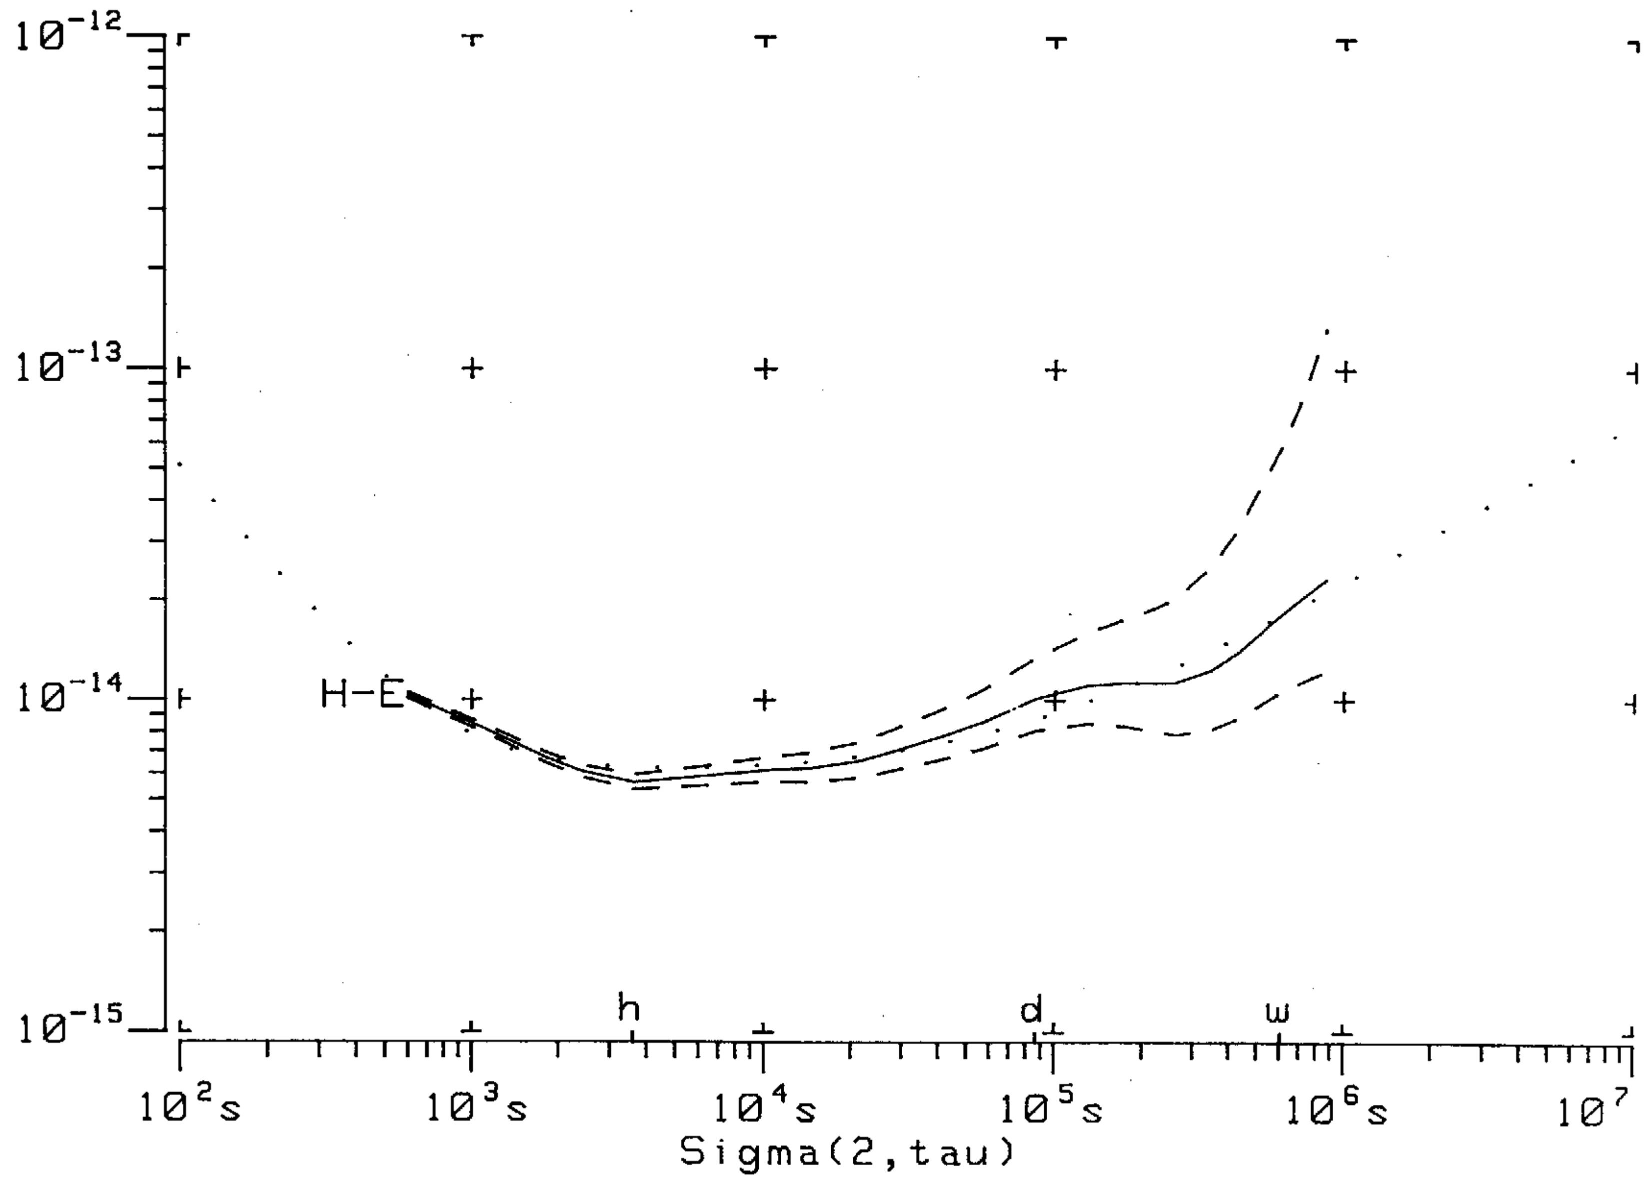

Figure 1. Allan deviation of H3 vs Ensemble for MJD=48795 to 48832 (37 days).

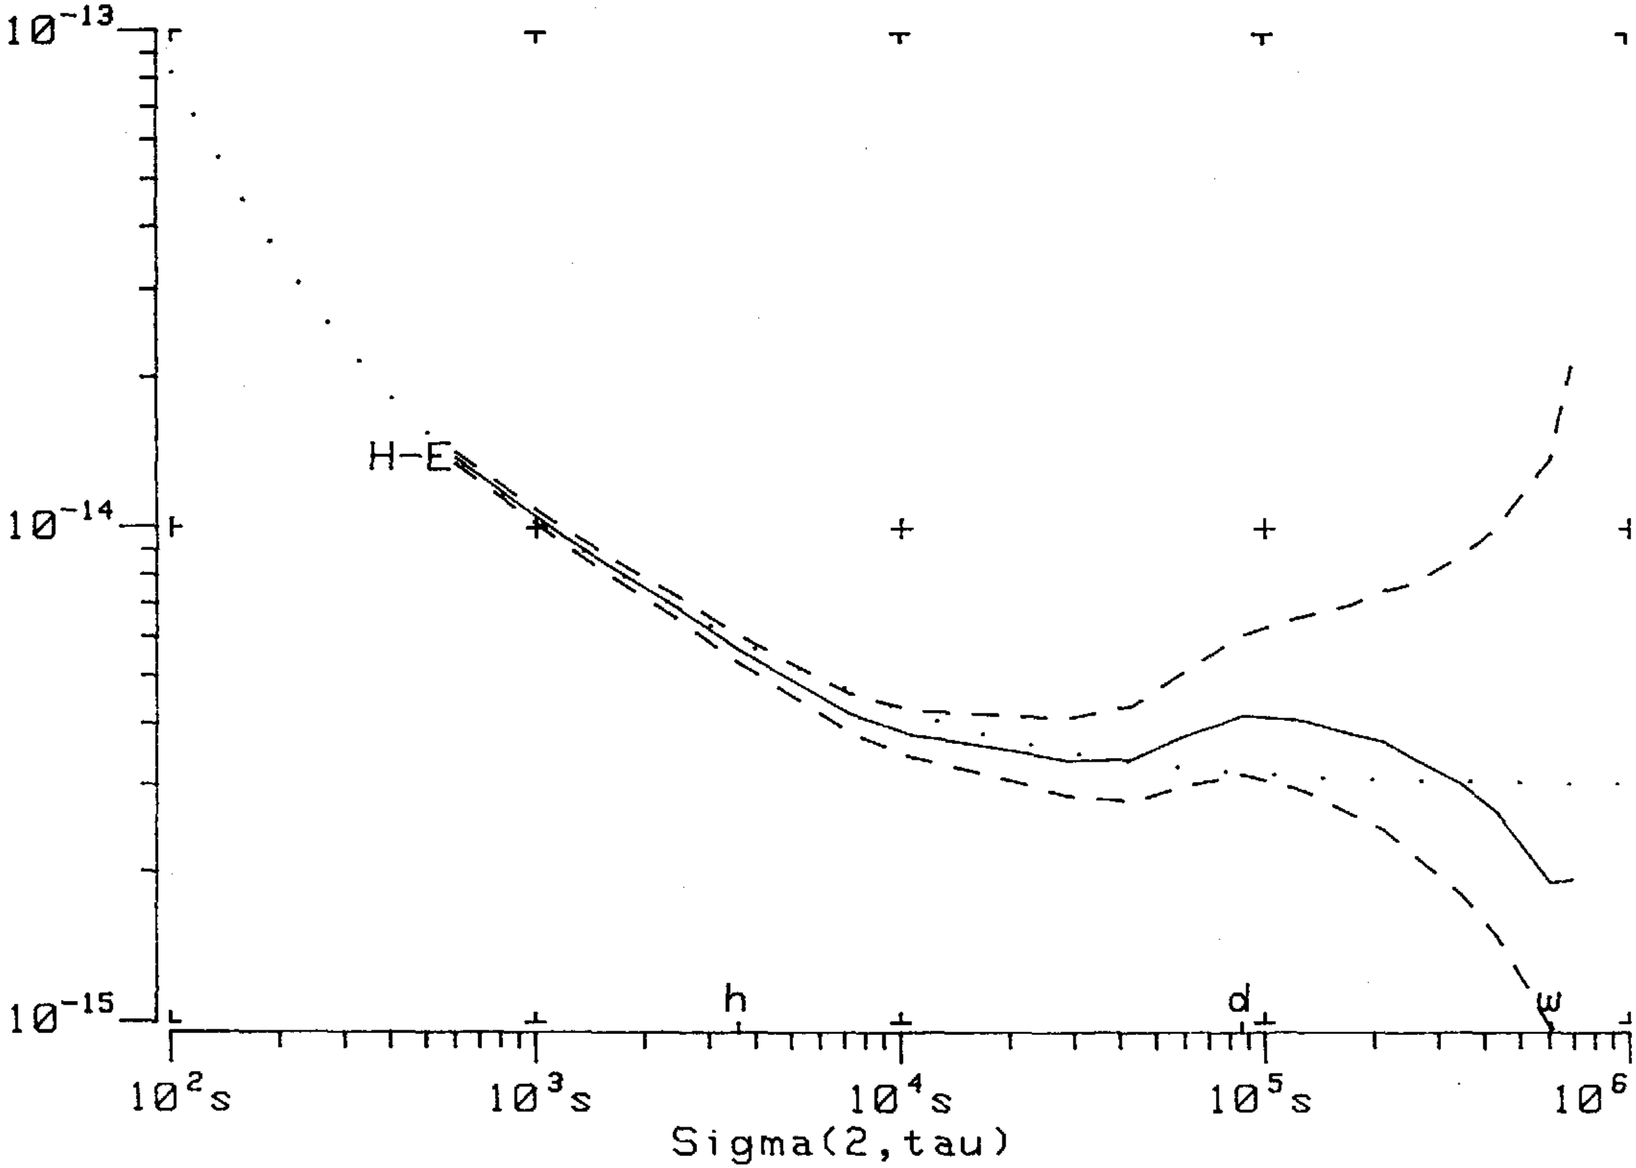

Figure 2. Allan deviation of H3 vs Ensemble for MJD=48895 to 48915 (20 days).

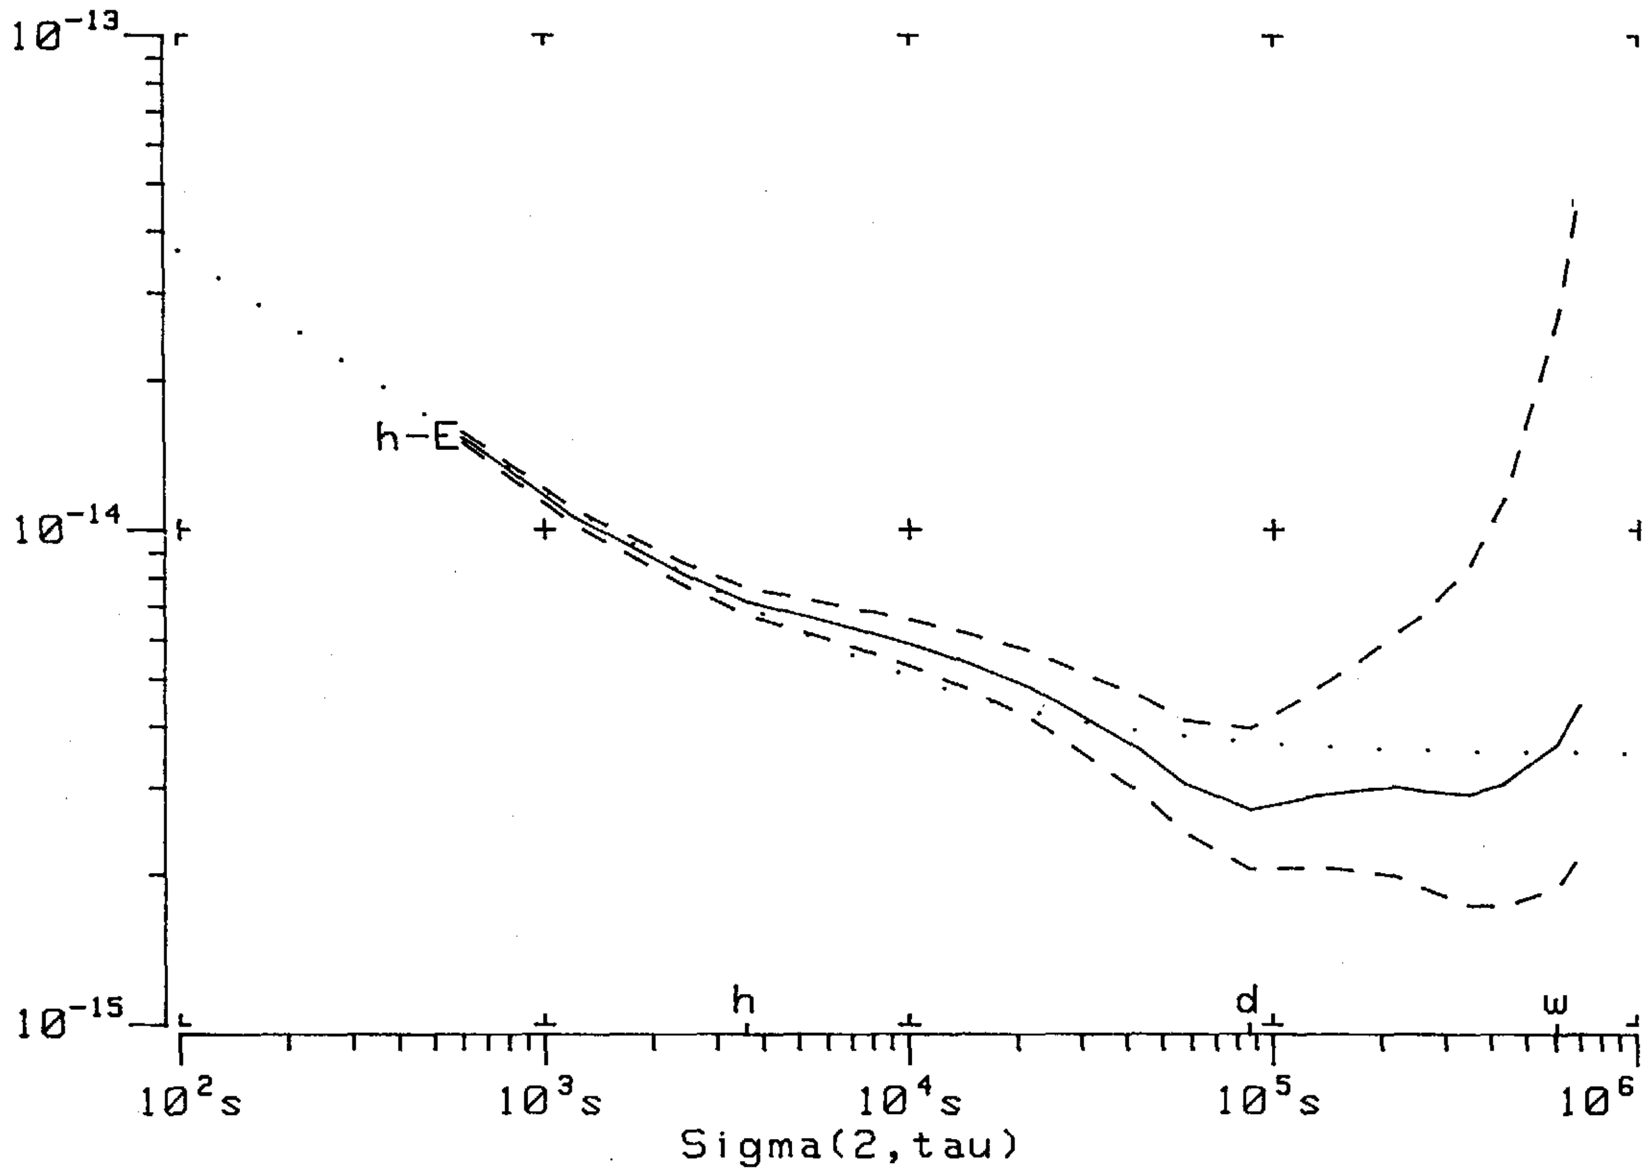

Figure 3. Allan deviation of H4 vs Ensemble for MJD=48895 to 48915 (20 days).

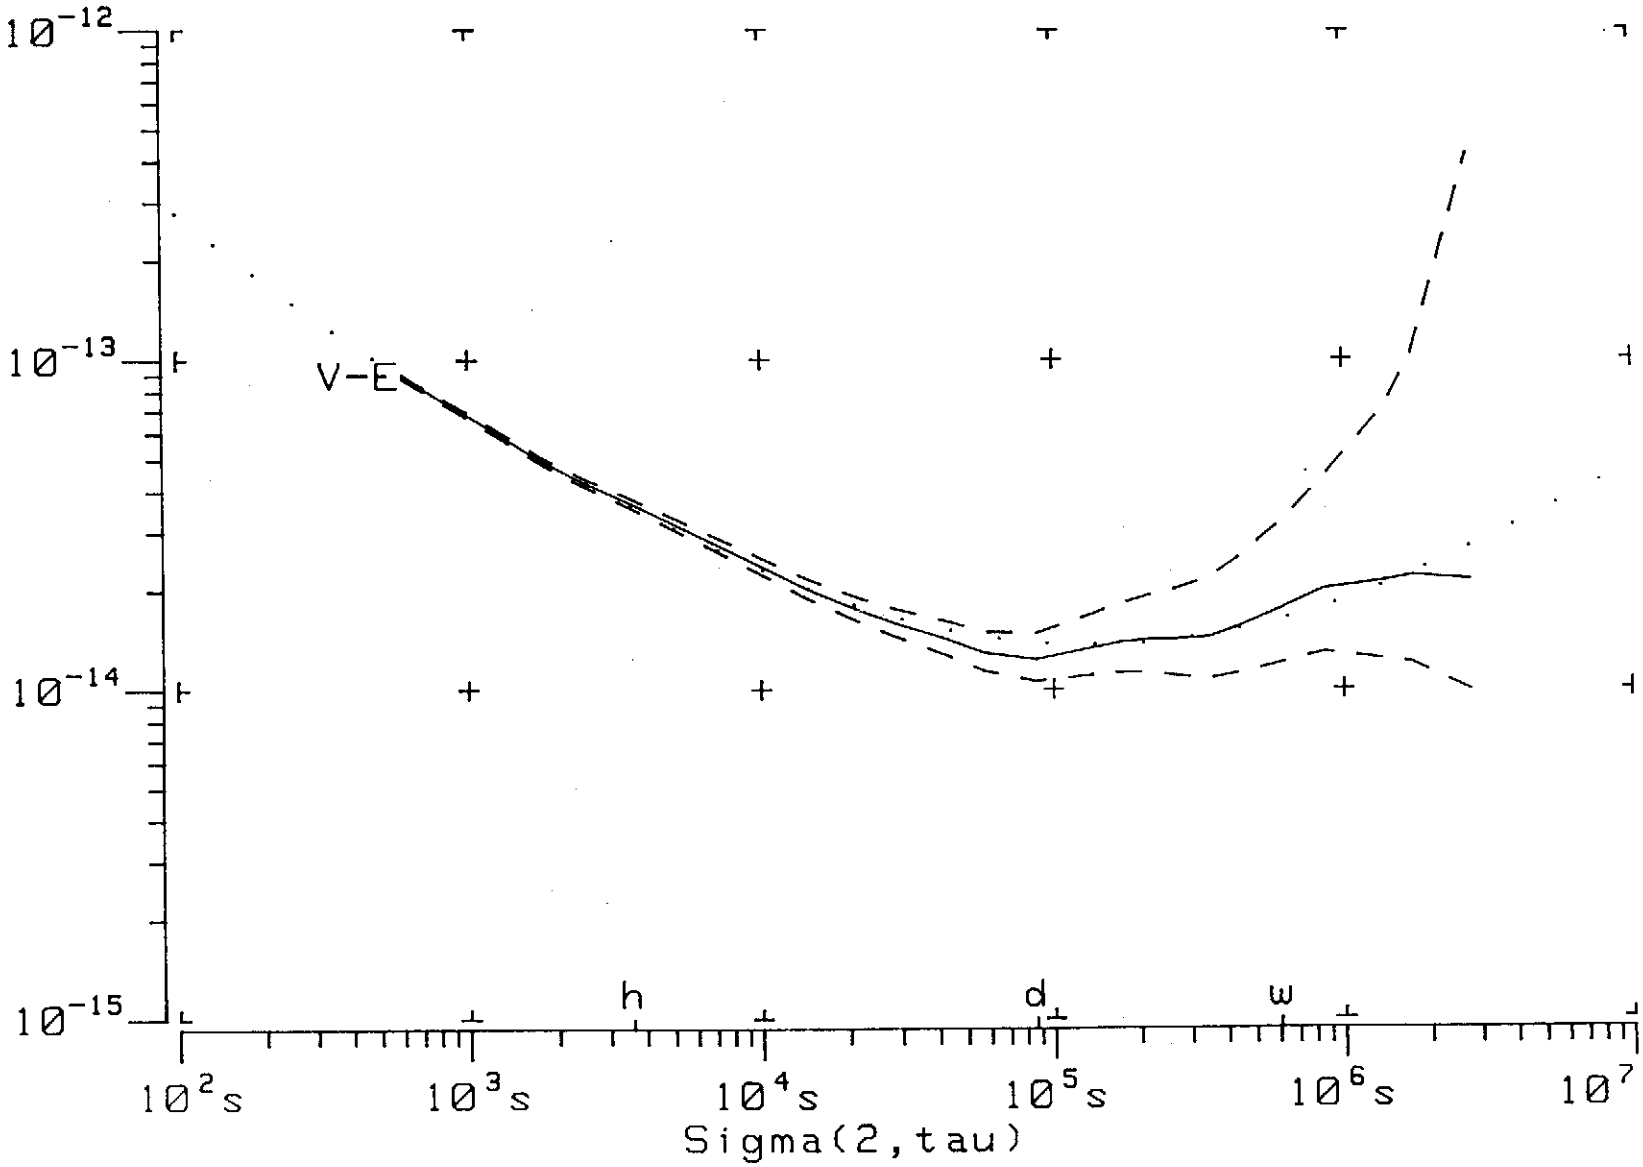

Figure 4. Allan deviation of Cs V vs Ensemble for MJD=48795 to 48865 (70 days).

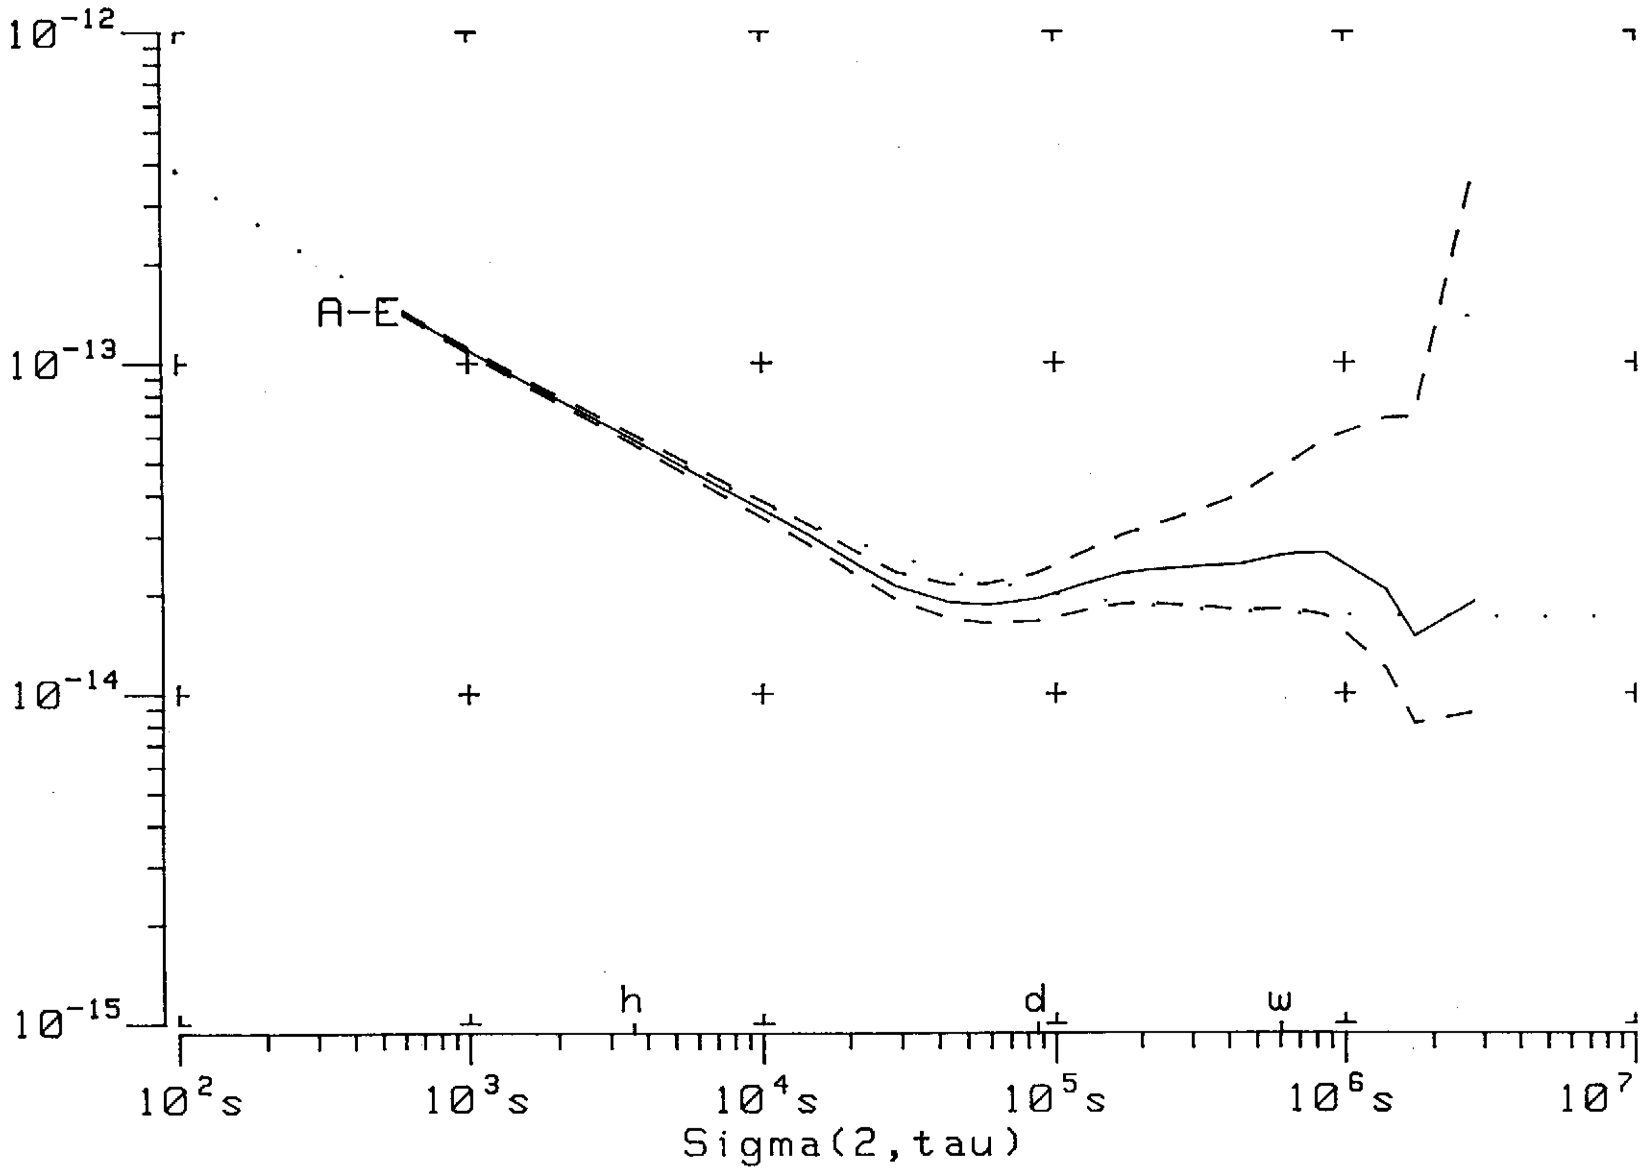

Figure 5. Allan deviation of Cs VI A vs Ensemble for MJD=48795 to 48865 (70 days).

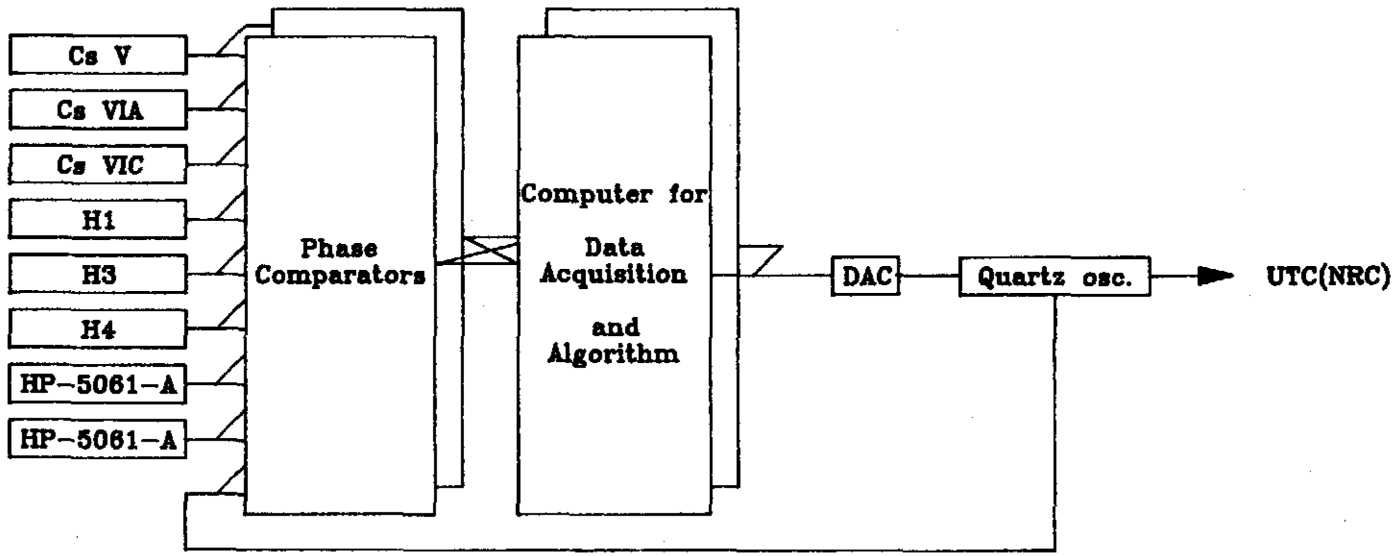

Figure 6. Time scale generation at NRC (project)