PRELIMINARY RESULTS FROM NPL'S CLOCK ENSEMBLE ALGORITHM USING HYDROGEN MASERS AND CAESIUM CLOCKS

Setnam L Shemar, John A Davis and Peter B Whibberley

National Physical Laboratory, Hampton Road, Teddington, Middlesex, TW11 0LW, UK
Email: setnam.shemar@npl.co.uk

ABSTRACT

In recent years the National Physical Laboratory (NPL) has been developing a Kalman filter-based atomic clock ensemble algorithm. The algorithm is being used to form a composite time scale which will ultimately be employed to improve the stability of the UK's national time scale UTC(NPL) and its robustness to hardware failures. The algorithm has been optimised for both short-term and long-term stability. Flicker Frequency Modulation (FFM) clock noise is modelled using a linear combination of integrated Markov noise processes. We present here recent results obtained for the composite time scale using data from an ensemble of three active hydrogen maser clocks and two caesium clocks. An initial assessment of the ensemble algorithm is carried out using a 150-day period of data and two subsets of the data. The results show that the output of the composite is consistent between different periods of clock data and also different lengths of clock data. Initial indications are that the composite is performing well.

INTRODUCTION

The National Physical Laboratory (NPL) maintains the UK's national time scale UTC(NPL). At present this is based on a Hydrogen-maser (H-maser) frequency standard, which is regularly checked for anomalies against other atomic standards including three additional active H-masers and two caesium clocks. In recent years NPL has been developing a Kalman filter-based clock ensemble algorithm as part of its long-term strategy for enhancing the performance and robustness of UTC(NPL). A description of the fundamental principles behind the algorithm together with initial results using an ensemble of three H-masers has been presented previously [1]. In this paper we describe additional results obtained from an ensemble of three Symmetricom Sigma-Tau active H-maser frequency standards and two Symmetricom 5071A caesium clocks.

ALGORITHM BACKGROUND

The development of Kalman filter-based algorithms for forming composite clocks has been described previously [2, 3, 4]. The fundamental principles employed in the algorithm used in this paper have also been described previously [1] and are outlined below.

The ensemble algorithm provides an estimate of the time-offset between each clock and the composite. It assumes that there is no measurement noise in the clock input data. The state vector comprises the parameters estimated by the algorithm. These include three classical components corresponding to the time offset, frequency offset and linear frequency drift offset from the composite for each clock. Flicker Frequency Modulation (FFM) is modelled as a linear combination of integrated Markov noise processes using additional state vector components. The composite has close to optimal performance in the presence of FFM [1].

For the nth iteration of the filter, and using the same notation as previously [1], it is assumed that the random process to be estimated can be modelled as follows

$$x_n = \Phi x_{n-1} + w_n \tag{1}$$

where x_n is the process state vector, Φ is the state propagation matrix and w_n is the process noise with zero expectation and covariance matrix Q and is independent of w_k for $k < n$. The relationship between the physical parameters being measured, y_n, and the state vector, x_n, is given as follows

$$y_n = Hx_n \tag{2}$$

where H represents the design matrix. It will be noted that there is no representation of the measurement error in (2), which is a consequence of the Kalman filter model assuming that the measurements are noiseless.

The well-known Kalman filter equations are employed as follows. A current estimate of the state vector is obtained by extrapolating $\hat{x}_{n-1}^{+}$ the previous estimate of the state vector using

$$\hat{x}_n^- = \Phi \hat{x}_{n-1}^+ . \tag{3}$$

Note that $\hat{x}_n^+$ represents the current state vector estimate using the current measurements y_n and $\hat{x}_n^-$ that not using y_n. Next, parameter covariance matrix extrapolation is carried out. This is normally carried out using the parameter covariance matrix of $\hat{x}_{n-1}^{+}$ but in this case

$$P_{\mathrm{R}n}^- = \Phi P_{\mathrm{R}(n-1)}^+ \Phi^{\mathrm{T}} + Q \tag{4}$$

where P_{Rn}^- , P_n^+ *and* P_{Rn}^+ are the reduced error parameter covariance matrices at the nth iteration of the filter [1, 3] , before and after the addition of the current measurements and after application of the covariance x-reduction process, respectively. The Kalman Gain $\mathrm{K_{Rn}}$ is determined as follows

$$K_{\mathrm{R}n} = P_{\mathrm{R}n}^- H^{\mathrm{T}} (H P_{\mathrm{R}n}^- H^{\mathrm{T}})^{-1} . \tag{5}$$

Next, the current measurements are incorporated to give a better current estimate of the state vector

$$\hat{x}_n^+ = \hat{x}_n^- + K_{\mathrm{R}n} (y_n - H\hat{x}_n^-) . \tag{6}$$

Next, the reduced error parameter covariance matrix is updated for the next iteration of the filter

$$P_n^+ = (I - K_{\mathrm{R}n} H) P_{\mathrm{R}n}^- \tag{7}$$

where I is an identity matrix with diagonal elements set to unity. The final equation below transforms P_n^+, using S the covariance x- reduction matrix, and is called 'covariance x-reduction' [1, 3]

$$P_{\mathrm{R}n}^+ = S P_n^+ S^{\mathrm{T}} . \tag{8}$$

The linear frequency drift present in an active H-maser may reach levels of $1\mathrm{x}10^{-19}\mathrm{s}^{-1}$. Such values will result in a time offset of order 30 ns at a prediction length of 10 days. Consequently, a significant non-zero value of a linear frequency drift term is used in the case of the active H-masers. In the case of the caesium clocks, the stochastic noise processes dominate relative to any linear frequency drift. Consequently, the linear drift parameter for each caesium clock is maintained very close to zero by using very small values for the relevant noise parameters in the covariance matrix Q and also having very small initial values for the relevant elements of the parameter covariance matrix. This approach is used to assist the estimation of the linear frequency drift of the H-masers relative to the caesium clocks. The model also assumes that the linear frequency drift of the active H-masers is not constant. Values very close to zero are used instead of zero where described above in order to avoid division by zero software computations. Although in future this parameter could be excluded for the caesium clocks altogether, at present it is in the software in order to allow maximum flexibility for future modifications and testing of the algorithm.

Separate noise parameters are used in the algorithm to model the noise of each clock type. The resulting Allan Deviation (ADEV) and Hadamard Deviation (HDEV) values at short averaging times are very much lower for the active H-masers than the caesium clocks. At long averaging times the resulting ADEV and HDEV values are closer, but still significantly lower in the case of the active H-masers. This effectively means that the active H-masers are given a much higher weight

than the caesium clocks at all averaging times. Different noise parameters can also be applied to clocks of the same type although this has not been done here.

INITIALISING THE ALGORITHM

The parameters that require initialisation are as follows.
(i) noise parameter estimates for each clock;
(ii) time offset, normalised frequency offset, linear frequency drift offset and integrated Markov noise parameters for each clock;
(iii) variance estimates of the time offset, normalised frequency offset, linear frequency drift and integrated Markov noise parameters used in the diagonal elements of the parameter covariance matrix, *P*, for each clock.

The 'noise parameters' correspond to the variances of the white noise processes underlying each clock noise type [5]. The performance of the algorithm will depend on the values of the noise parameters set and it is important that these reliably match those of the actual clock noise processes. The algorithm uses a preliminary set of these to determine ADEV and HDEV values for each type of clock, which provides the operator with a useful check on the suitability of the noise parameter choice. The noise parameters are then refined using separate software such that they better match the ADEV and HDEV characteristics of the manufacturer's clock specification. These improved values are then used for initialisation.

The frequency drift offset for the caesium clocks is always initialised very close to zero. Good initialisation of the other individual clock offsets is vital to minimize the time required for the algorithm to achieve steady state. For the results presented below, the linear frequency drift of the active H-masers have also been initialised very close to zero. Given a good initialisation the algorithm may achieve optimal performance within a few hours or less. This may even be achieved immediately given a good enough initialisation. Most important is initialising the linear frequency drift estimates for individual active H-masers. These parameters require a long period of input data before the algorithm is able to produce estimates with low uncertainties. Consequently, the performance of the algorithm in terms of ADEV and HDEV stability of the composite may be degraded during this period.

The diagonal elements of the parameter covariance matrix for each type of clock should be set to levels that correctly reflect the uncertainty of the initial calibration. These correspond to the uncertainties in the initial values of the state vector components.

INPUT CLOCK DATA

Data from three of NPL's H-maser frequency standards and two caesium clocks have been used to carry out an initial assessment of the ensemble algorithm. The data are derived from measurements between the 1 Pulse Per Second (1 PPS) output of each clock and that of a reference clock. The measurements are taken once per hour and correspond to the mean of a series of 10 measurements taken at one second intervals for each clock using a Stanford Research Systems model SR620 counter timer. The active H-maser that provides UTC(NPL) (also referred to as Clock 2 in the results presented below) has been selected as the reference clock and the data correspond to counter timer measurements of (UTC(NPL) – C_I) where C_I is an individual clock.

The occasional frequency steers applied to UTC(NPL) have been taken into account within the ensemble algorithm. This is achieved by entering the size of a steer in terms of the fractional frequency change (in units of ns/day) directly into the frequency offset component of the state vector of the clock at the relevant MJD of the steer.

RESULTS

Here we present the ensemble algorithm output for a 150 day period of clock measurements spanning MJDs of 54539 to 54689. For purposes of checking the consistency of the algorithm's output for different samples of input clock data we also present the output for two non-overlapping 60 day subsets of the 150 day period separated by 10 days. The two subsets correspond to periods with MJDs from 54539 to 54599 and from 54609 to 54669. These two specific periods were chosen because there are two frequency steers applied to UTC(NPL) within the 150 day period at MJDs 54605.69 and 54670.73 (corresponding to fractional frequency steers of -0.998 ns/day and 0.852 ns/day respectively) which do

not fall in either of the 60 day periods. This means that no correction is required for frequency steers, enabling a simpler analysis for the 60 day periods.

Fig. 1 shows the complete 150 day period of input data in terms of time-offset measurements between a reference clock and each individual clock. Clocks 1, 2 and 3 are active H-masers, whilst Clocks 4 and 5 are caesium clocks. The reference clock is Clock 2 which also corresponds to UTC(NPL) with a fixed offset. Consequently, Fig. 1 shows measurements of (UTC(NPL) – C_i) versus MJD, where C_i corresponds to the *i*th individual clock in the ensemble. It can also be seen that there is an arbitrary offset in the Clock 2 measurements resulting from the measurement system. There is a short period between MJDs 54676.292 and 54676.458 corresponding to five one hour-interval measurements where data for Clock 3 are missing. Although the algorithm can be configured to continue generating a composite across such a period, this is not the subject of this paper. Instead an approximate linear interpolation is carried out using the data points immediately before and after the data gap and used to simulate data points where missing. Due to the short period over which this has been carried out, it is not expected to have any significant effect on the results presented. Fig. 1 also highlights the two 60-day subsets of data. It is notable that there is no obvious visible linear frequency drift between any individual clock and Clock 2.

Fig. 2 shows the same data as in Fig. 1 but in the form of a series of three plots corresponding to (a) the first 60 day period corresponding to MJDs between 54539 and 54599, (b) the second 60 day period corresponding to MJDs between 54609 and 54669 and (c) the complete 150 day period corresponding to MJDs between 54539 and 54699. The y-axes for all three plots in the figure are the same as shown in (a) and this also applies to all subsequent figures. The relative clock stabilities of the (UTC(NPL) – C_i) measurements in terms of ADEV and HDEV for the three periods of data are shown in Figs. 3 and 4 respectively. The very low ADEV and HDEV results for Clock 2 are a consequence of this clock being used as the reference clock, resulting in the ADEV in this case showing only the noise in the measurement system. It can be seen from Figs. 3 and 4 that the ADEV and HDEV plots are similar to each other for corresponding clock pairs at the longer averaging times. As HDEV is a statistic that is insensitive to linear frequency drift, this indicates that there is little or no significant linear frequency drift between the three active H-masers observed with averaging times of up to 10 days.

The output of the ensemble algorithm is given in Fig. 5 which shows time offset measurements between each individual clock and the composite. It is notable that there is no visible linear frequency drift between any individual clock and the composite. Figs. 6 and 7 show the ADEV and HDEV results respectively of the time offset measurements shown in Fig. 5. It can be seen that the ADEV and HDEV plots for Clocks and 1 and 3 show stabilities relative to the composite that are consistent with those relative to Clock 2 shown in Figs. 3 and 4. However, in the case of Clock 2 the stabilities relative to the composite for the two 60 day periods shown in Figs. 5 and 6 are significantly better than for those relative to Clock 2 shown in Figs. 3 and 4. This indicates that the composite is showing a good performance in terms of stability. Furthermore, it can be seen that the ADEV and HDEV values for Clock 2 are considerably lower for the two 60 day periods than for the 150 day period. This is simply a consequence of frequency steers having been applied to Clock 2 within the 150 day period but not in the two 60 day periods. It is also apparent that the ADEV and HDEV plots are similar to each other at longer averaging times for the three active H-masers suggesting that there is no significant linear frequency drift between them and the composite.

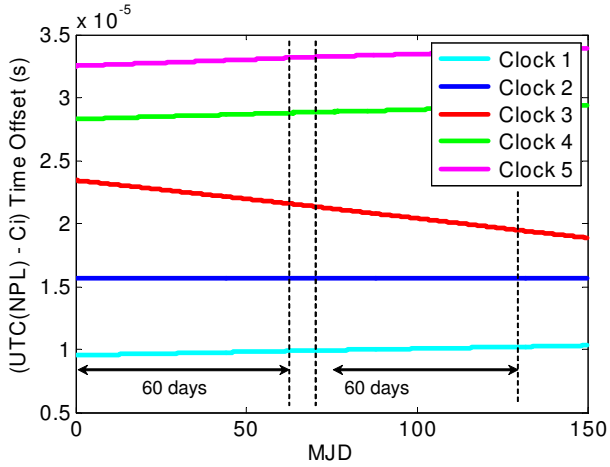

Fig. 1. Plot showing the individual clock differences from Clock 2 (also referred to as UTC(NPL)).

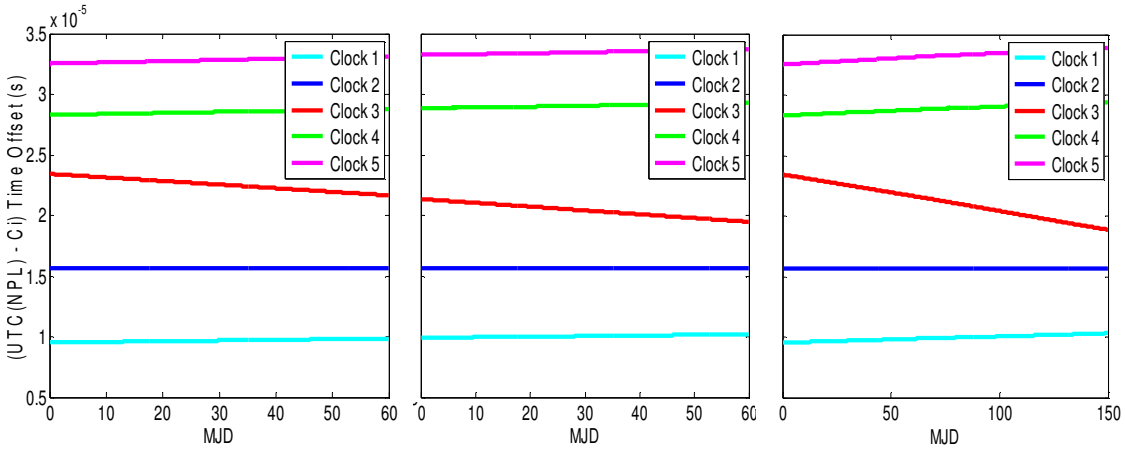

(a) (b) (c)

Fig. 2. Plots showing the individual clock differences from Clock 2 (also referred to as UTC(NPL)) for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

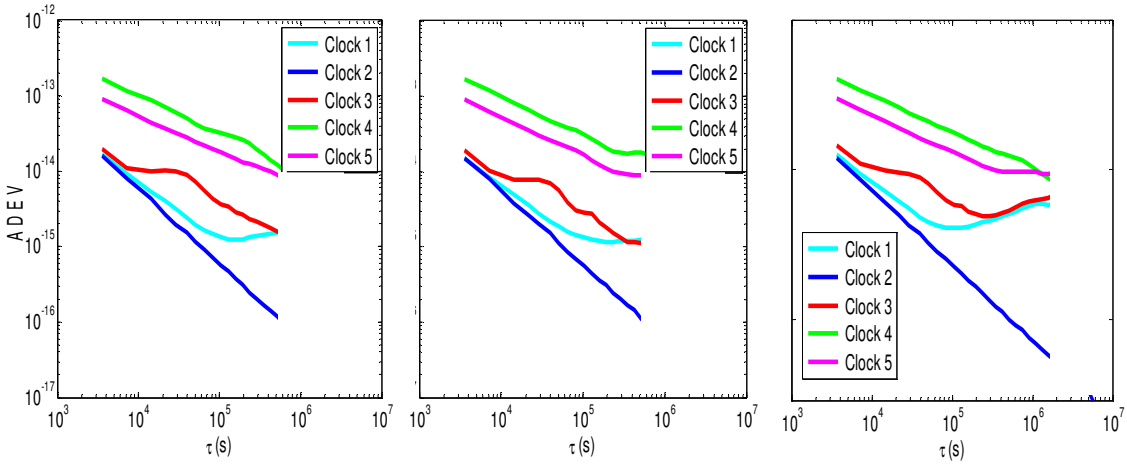

(a) (b) (c)

Fig. 3. Plots showing the ADEV of individual clock differences from Clock 2 (also referred to as UTC(NPL)) for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

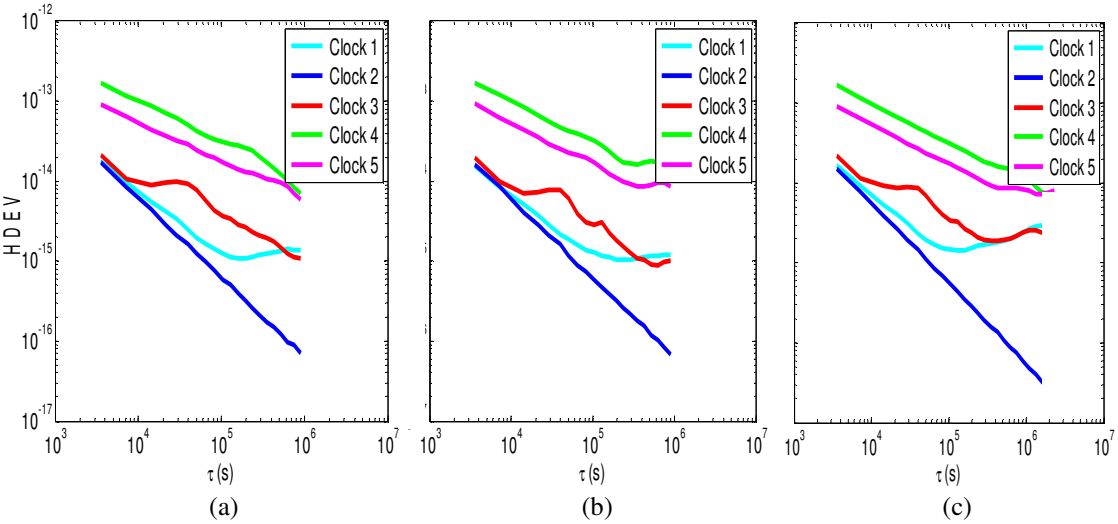

(a) (b) (c)

Fig. 4. Plots showing the HDEV of individual clock differences from Clock 2 (also referred to as UTC(NPL)) for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

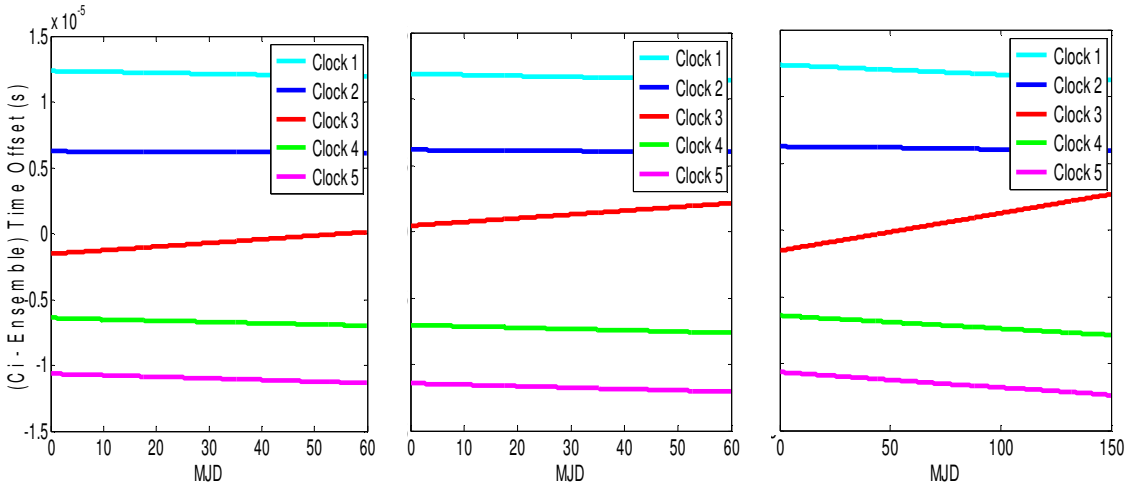

(a) (b) (c)

Fig. 5. Plots showing the individual clock differences between each individual clock and the composite for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

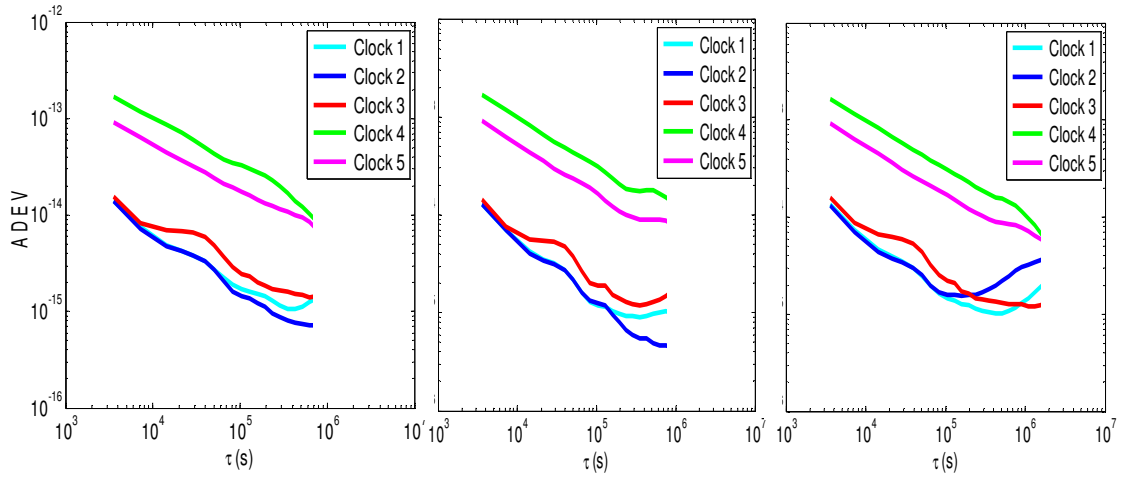

(a) (b) (c)

Fig. 6. Plots showing the ADEV of individual clock differences from the composite for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

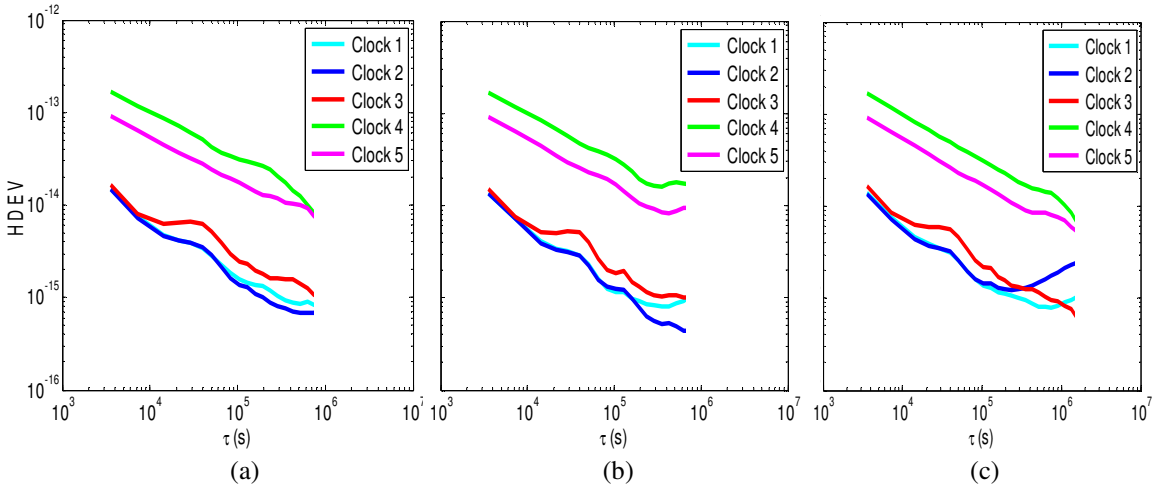

(a) (b) (c)

Fig. 7. Plots showing the HDEV of individual clock differences from the composite for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

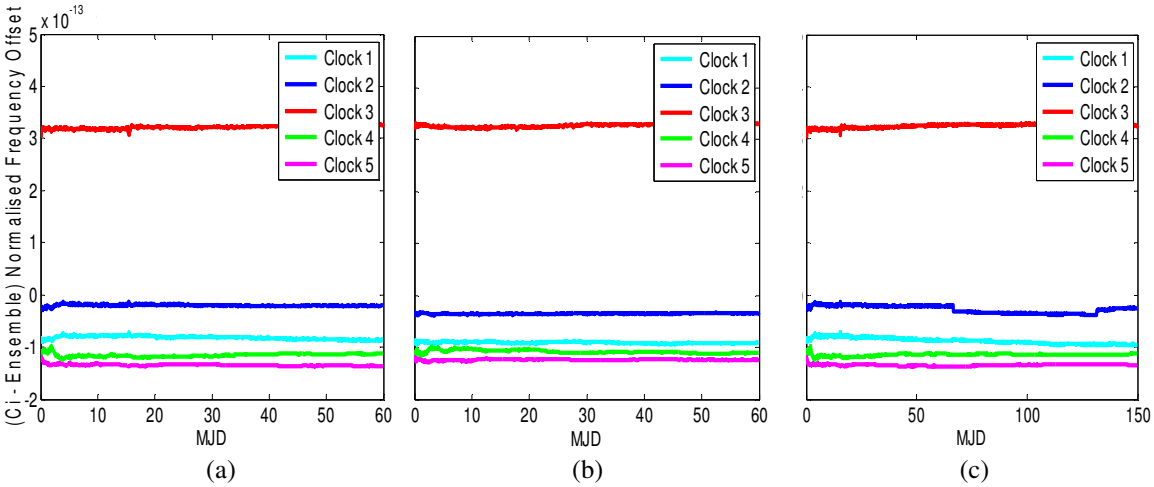

Fig. 8. Plots showing the normalised frequency offset between individual clocks and the composite for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

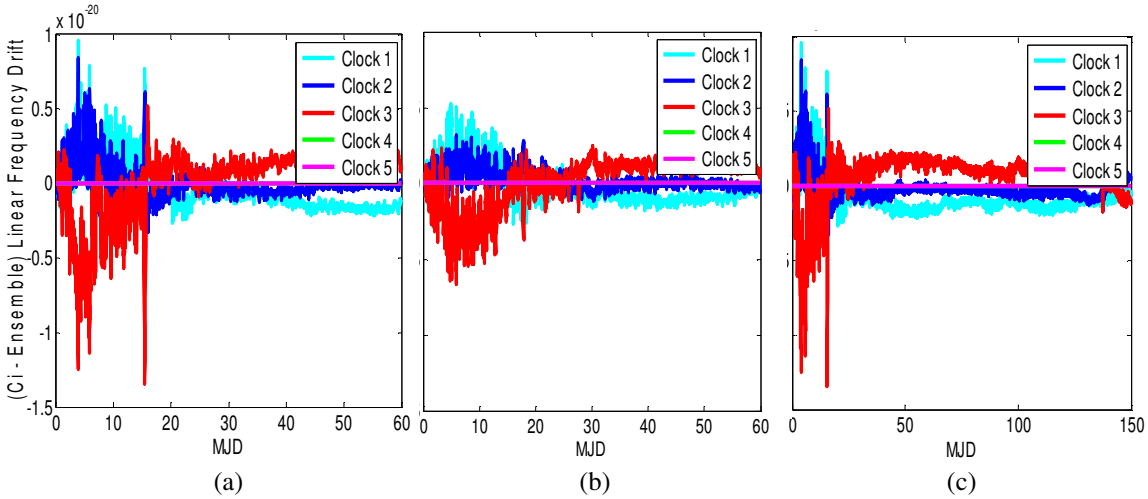

Fig. 9. Plots showing the linear frequency drift offset between individual clocks and the composite for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period.

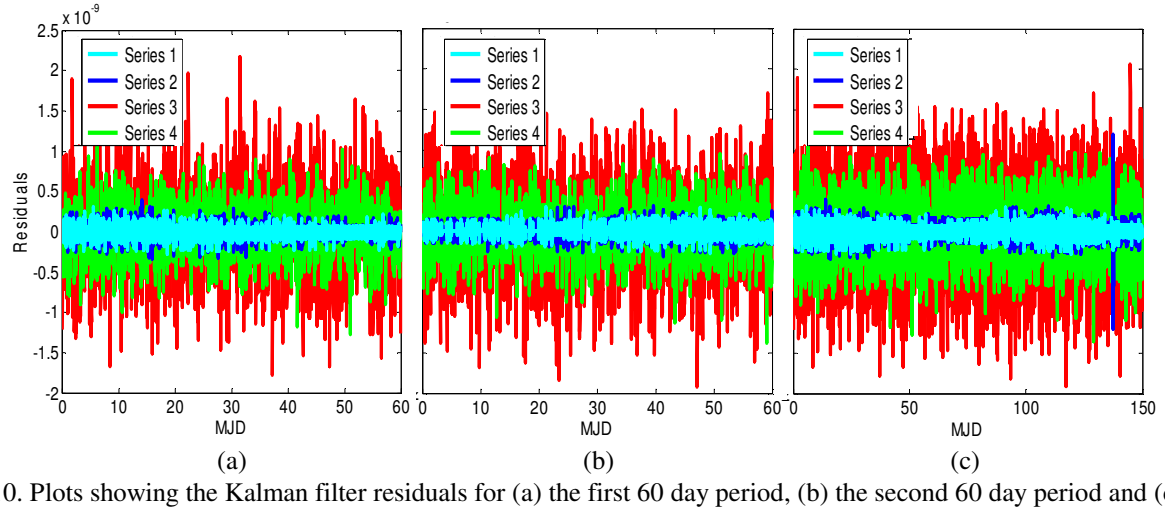

Fig. 10. Plots showing the Kalman filter residuals for (a) the first 60 day period, (b) the second 60 day period and (c) the 150 day period. Series 1, 2, 3 and 4 correspond to the residuals for the time offsets from Clock 2 of each of Clocks 1, 3, 4 and 5 respectively.

Fig. 8 shows the normalised frequency offset between each individual clock and the composite. It can be seen that the frequency offsets are stable and consistent between the two 60 day periods. In the case of the 150 day period, the frequency steers applied to UTC(NPL) are clearly visible in the data for Clock 2.

Fig. 9 shows the linear frequency drift offset between each clock and the composite. Estimation of the linear frequency drift offsets of the H-masers is assisted by setting the linear frequency drift offsets of the caesium clocks very close to zero. Fig. 9 shows that the linear frequency drift offsets of the caesium clocks have been set very close to zero within the algorithm.

Fig. 10 shows the Kalman filter residuals within the algorithm. These correspond, in each iteration of the filter, to the difference between the current measurement and the filter's prediction using the previous measurement. The very brief increase in residuals for Clock 3 towards the end of the 150 day period seen in (c) is an artefact due to the approximate linear interpolation used to simulate data points across the period of missing data for this clock between MJDs 54676.292 and 54676.458. Apart from this, it is apparent that the residuals for each clock show no visible inconsistencies between the three periods of data, or visible biases between clocks.

CONCLUSIONS

We have presented a preliminary assessment of a Kalman filter-based clock ensemble algorithm using input data from three active H-masers and two caesium clocks. The algorithm is being used to form a composite time scale which will ultimately be employed to improve the stability of the UK's national time scale UTC(NPL) and its robustness to hardware failures.

Estimation of the linear frequency drift offsets of the H-masers is assisted by initialising and maintaining the linear frequency drift offsets of the caesium clocks very close to zero. The model also assumes that the linear frequency drift of the active H-masers is not constant.

It is found that the composite time scale gives consistent results when using different periods of input clock data and also different lengths of input data. In the medium term of up to 10 days averaging time the composite time scale does not show any significant linear frequency drift. Initial indications are that the composite is performing well.

ACKNOWLEDGEMENTS

This work was supported by the UK Department of Business, Innovation and Skills as part of the National Measurement System Physical Metrology programme.

REFERENCES

[1] J. A. Davis, C. A. Greenhall and P. W. Stacey, "A Kalman filter clock algorithm for use in the presence of flicker frequency modulation noise", *Metrologia*, vol. 42, 2005.

[2] K. R. Brown, "The theory of the GPS composite clock", *Proceedings of ION GPS-91*, pp 223-41, 1991.

[3] C. A. Greenhall, "Forming stable time scales from the Jones-Tryon Kalman filter", *Metrologia*, vol. 40, S335-341, 2003.

[4] L. Galleani and P Tavella, "On the use of the Kalman filter in timescales", *Metrologia*, vol. 40, S326-S334, 2003.

[5] P. M. Harris, J. A. Davis, M. G. Cox and S. L. Shemar, "Least squares analysis of time series data and its application to two-way satellite time and frequency transfer measurements", *Metrologia*, vol. 40, S342-S347, 2003.