

KLTS: A rigorous method to compute the confidence intervals for the Three-Cornered Hat and for Gros Lambert Covariance

Éric Lantz^{1*}, Claudio E. Calosso², Enrico Rubiola^{3,2}, Vincent Giordano³, Christophe Fluhr⁴, Benoît Dubois⁴ and François Vernotte^{3*}

Abstract—The three-cornered hat / Gros Lambert Covariance methods are widely used to estimate the stability of each individual clock in a set of three, but no method gives reliable confidence intervals for large integration times.

We propose a new KLTS (Karhunen-Loève Transform using Sufficient statistics) method which uses these estimators to take into account the statistics of all the measurements between the pairs of clocks in a Bayesian way. The resulting Cumulative Density Function (CDF) yields confidence intervals for each clock AVAR. This CDF provides also a stability estimator which is always positive.

Checked by massive Monte-Carlo simulations, KLTS proves to be perfectly reliable even for one degree of freedom. An example of experimental measurement is given.

Keywords— Clock stability; Allan variance; three-cornered hat; covariances; confidence interval; Bayesian analysis

I. INTRODUCTION

Although the three-cornered hat [1] and the Gros Lambert Covariance (GCov) [2] methods are widely used to measure the stability of each individual clock in a set of three, the only methods which exist to compute error bars are limited to the smallest integrations times, i.e. when the number of Equivalent Degrees of Freedom (EDF) is high [3], [4], [5]. However, there is no reliable method to assess confidence intervals over the estimates if their number of EDF is low, typically 5 or below. However, since this case occurs for the largest integration times, it is an important issue for all applications dealing with long term stability (e.g. time keeping).

Likewise, another problem frequently arises using the three-cornered hat or the GCov method: negative variance estimates can occur. Although this issue was already considered by Premoli and Tavella [6], it would be useful to get a method which could provide simultaneously a positive estimate as well as its confidence interval.

In a previous paper, we performed a first Bayesian attempt to estimate confidence intervals from the statistics of the three-cornered estimates but we observed that this method was only valid beyond 5 EDF [5]. We propose here a new method, also based on Bayesian inversion, which uses the statistics of the data themselves, instead of the statistics of the estimates.

¹Département d’Optique P.M. Duffieux, FEMTO-ST, UMR 6174 CNRS-Université Bourgogne Franche-Comté, France

²Physics Metrology Division, Istituto Nazionale di Ricerca Metrologica INRiM, Torino, Italy

³Department of Time and Frequency, FEMTO-ST, Observatory THETA, UMR 6174 CNRS-Université Bourgogne Franche-Comté, France

⁴FEMTO Engineering, Besançon, France

*Corresponding authors: eric.lantz@univ-fcomte.fr, francois.vernotte@obs-besancon.fr




Fig. 1. Layout of the clock measuring device.

Nevertheless, these statistics of the data can be computed from the values of the estimates, since these estimates form a “sufficient statistics” [7] for the variance estimation. The resulting Cumulative Density Functions (CDF) yield the lower and upper bounds of the 95 % confidence interval. On the other hand, the CDF at 50 % provides a useful stability estimator of the stability of each clock, i.e. the median value, which has the advantage of always being positive.

The performances of this method have been checked by using massive Monte-Carlo simulations. The principle of these simulations is described in this paper and the comparisons with the theoretical confidence intervals given by our new method are discussed. Finally, this method is applied to the measurement of 3 Cryogenic Sapphire Oscillators (CSO).

II. STATEMENT OF THE PROBLEM

A. Clock comparisons

Let us consider 3 independent clocks: A , B and C . It is possible to compare these clocks by pairs by using 3 Time Interval Counter (TIC), as shown in Figure 1, and to estimate the corresponding Allan variances (AVAR). Let us denote \bar{y}_{ABk} the k^{th} frequency deviation sample between A and B integrated without dead time over a duration τ , $\bar{y}_{ABk} = \frac{1}{\tau} \int_{t_k}^{t_k+\tau} y(t) dt$, and $z_{ABk} = (\bar{y}_{ABk+1} - \bar{y}_{ABk}) / \sqrt{2}$. For clocks A and B , the AVAR is

$$\sigma_{AB}^2(\tau) = \mathbb{E} [z_{ABk}^2] \quad (1)$$

where $\mathbb{E}[\cdot]$ stands for the mathematical expectation of the quantity between the brackets. For the sake of concision of the notations, we will drop the dependence on τ in the following. Since $\sigma_{AB}^2 = \sigma_A^2 + \sigma_B^2$, the three-cornered hat is based on the following property:

$$\text{TCH}_A = \frac{1}{2} (\sigma_{AB}^2 - \sigma_{BC}^2 + \sigma_{CA}^2) = \sigma_A^2. \quad (2)$$

On the other hand, the Gros Lambert covariance is based on this other property:

$$\text{GCov}_A = \mathbb{E}[\mathbf{z}_{ABk} \cdot \mathbf{z}_{ACk}] = \sigma_A^2. \quad (3)$$

B. Measurement noise influence

Let us consider the layout of Figure 1:

- the clocks B and C are connected to the counter 1 which is affected by a measurement noise $\{\mathbf{z}_{N1k}\}$
- the clocks C and A are connected to the counter 2 which is affected by a measurement noise $\{\mathbf{z}_{N2k}\}$
- the clocks A and B are connected to the counter 3 which is affected by a measurement noise $\{\mathbf{z}_{N3k}\}$.

We assume that these 3 measurement noises are uncorrelated.

The measurement given by the counter 3 is then the sum of the clock noises plus the corresponding measurement noise: $\mathbf{z}_{3k} = \mathbf{z}_{ABk} + \mathbf{z}_{N3k}$. Its variance is

$$\sigma_3^2 = \mathbb{E}[(\mathbf{z}_{ABk} + \mathbf{z}_{N3k})^2] = \sigma_{AB}^2 + \sigma_{N3}^2 = \sigma_A^2 + \sigma_B^2 + \sigma_{N3}^2. \quad (4)$$

In the following, we will call \mathbf{z}_{1k} , \mathbf{z}_{2k} and \mathbf{z}_{3k} the **measurements**.

The three-cornered hat becomes $\text{TCH}_A = \frac{1}{2} (\sigma_3^2 - \sigma_1^2 + \sigma_2^2) = \sigma_A^2 + \frac{1}{2} (\sigma_{N3}^2 - \sigma_{N1}^2 + \sigma_{N2}^2)$. Let us assume that the 3 counter measurement noises are of the same level: $\sigma_{N1}^2 \approx \sigma_{N2}^2 \approx \sigma_{N3}^2 = \sigma_N^2$. Therefore $\text{TCH}_A = \sigma_A^2 + \frac{1}{2} \sigma_N^2$.

By contrast, the Gros Lambert covariance remains, because the measurement noises are uncorrelated:

$$\text{GCov}_A = \mathbb{E}[(\mathbf{z}_{ABk} + \mathbf{z}_{N3k}) \cdot (\mathbf{z}_{ACk} + \mathbf{z}_{N2k})] = \sigma_A^2. \quad (5)$$

Therefore, the only difference between these two approaches concerns the measurement noise due to the counters, since the expectation of the GCov estimates is not sensitive to this noise [4]. Note however that the variance of the GCov estimates does include a measurement noise term.

C. Model parameters and estimates

The previously defined quantities σ_A^2 , σ_{AB}^2 , TCH_A , GCov_A , σ_1^2 , ... are unknown real values. However, they can be estimated by estimates which are random variables. We can then define the following estimates:

$$\begin{aligned} \hat{\sigma}_3^2 &= \frac{1}{M} \sum_{k=1}^M \mathbf{z}_{3k}^2 = \frac{1}{M} \sum_{k=1}^M (\mathbf{z}_{ABk} + \mathbf{z}_{N3k})^2 \\ &= \hat{\sigma}_A^2 + \hat{\sigma}_B^2 + \hat{\sigma}_N^2 \end{aligned} \quad (6)$$

where the hat ($\hat{\cdot}$) stands for *estimate*, M is the number of available consecutive \mathbf{z}_{ABk} . Let us call the $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$ and $\hat{\sigma}_3^2$ the **elementary estimates**.

Similarly, $\widehat{\text{TCH}}_A = \hat{\sigma}_A^2 + \frac{1}{2} \hat{\sigma}_N^2$ and

$$\begin{aligned} \widehat{\text{GCov}}_A &= \frac{1}{M} \sum_{k=1}^M \mathbf{z}_{3k} \mathbf{z}_{2k} \\ &= \frac{1}{M} \sum_{k=1}^M (\mathbf{z}_{ABk} + \mathbf{z}_{N3k}) (\mathbf{z}_{ACk} + \mathbf{z}_{N2k}) \\ &= \hat{\sigma}_A^2. \end{aligned} \quad (7)$$

Let us call the $\hat{\sigma}_A^2$, $\hat{\sigma}_B^2$ and $\hat{\sigma}_C^2$ the **final estimates**.

Meanwhile, let us call the σ_A^2 , σ_B^2 and σ_C^2 the **model parameters**.

The aim of this paper consists in calculating a confidence interval over each model parameter σ_A^2 , σ_B^2 and σ_C^2 , from the final and elementary estimates.

D. Estimation of the measurement noise by the closure relationship

The closure relationship is obtained by computing the sum of the measurements of all counters for a given k :

$$\begin{aligned} \mathbf{z}_{cls,k} &= \mathbf{z}_{1k} + \mathbf{z}_{2k} + \mathbf{z}_{3k} \\ &= \mathbf{z}_{BCk} + \mathbf{z}_{N1k} + \mathbf{z}_{CAk} + \mathbf{z}_{N2k} + \mathbf{z}_{ABk} + \mathbf{z}_{N3k} \\ &= \mathbf{z}_{N1k} + \mathbf{z}_{N2k} + \mathbf{z}_{N3k}. \end{aligned} \quad (8)$$

since $\mathbf{z}_{ABk} = \mathbf{z}_{Bk} - \mathbf{z}_{Ak}$. These 3 measurement noises being uncorrelated and of equal level, the variance of the closure is:

$$\sigma_{cls}^2 = \mathbb{E}[(\mathbf{z}_{N1k} + \mathbf{z}_{N2k} + \mathbf{z}_{N3k})^2] = 3\sigma_N^2.$$

This gives an efficient way to estimate the variance of the measurement noise:

$$\hat{\sigma}_N^2 = \frac{1}{3M} \sum_{k=1}^M (\mathbf{z}_{N1k} + \mathbf{z}_{N2k} + \mathbf{z}_{N3k})^2 = \frac{1}{3} \hat{\sigma}_{cls}^2. \quad (9)$$

E. Bayesian inference

In Bayesian analysis, we have to consider the model parameters $\Theta = (\theta_1, \dots, \theta_m)^T$ which have, in the model world, m definite but unknown values, and the measurements $\vec{X} = (x_1, \dots, x_n)^T$ which are n random variables. In our case, the parameters are the 3 true σ_A^2 , σ_B^2 and σ_C^2 AVAR values of the 3 clocks and the measurements are either the elementary estimates $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$ and $\hat{\sigma}_3^2$ (see the previous method described in [5]) or directly the $3M$ measurements $\{\mathbf{z}_{ABk}, \mathbf{z}_{BCk}, \mathbf{z}_{ACk}\}$. In the present method, we will compute the Probability Density Function (PDF) of these $3M$ measurements, by using only the measured values of the final and elementary estimates, because these estimates form a sufficient statistics for the measurements [7], [8]: the precise knowledge of the set of measurements does not bring any new information beyond the estimates.

While using the same estimates as [5], the present method computes only the gaussian PDF of the measurements themselves, while [5] lays on the approximation of the PDF of the estimators by a Gaussian law, inducing errors for a low number of EDF.

The Bayesian inversion lays on the distinction between two issues:

- the **direct problem**, which consists, in the model world, in calculating the Probability Density Function (PDF) of the estimates knowing the model parameters $p(\vec{X}|\vec{\Theta})$;
- the **inverse problem**, which consists, in the experimenter world, in calculating the PDF of the model parameters knowing the estimates $p(\vec{\Theta}|\vec{X})$. This is the most precise knowledge we can gain on these parameters after a measurement.

The direct problem has been solved in [5] and the inverse problem may be solved thanks to the Bayes theorem:

$$\begin{cases} p(\vec{\Theta}|\vec{X}) \propto \pi(\vec{\Theta}) \cdot p(\vec{X}|\vec{\Theta}) \\ \int p(\vec{\Theta}|\vec{X}) d\vec{X} = 1 \end{cases} \quad (10)$$

where $\pi(\vec{\Theta})$, called the prior, is the *a priori* probability of the parameter $\vec{\Theta}$ before any measurement.

III. THE KLTS METHOD

In the following, in order to distinguish our previous method described in [5] and the present method, we will call the former the KLTG method, for ‘‘Karhunen-Loève Transform with Gaussian approximation’’ method, and the latter the KLTS method, for ‘‘Karhunen-Loève Transform using Sufficient statistics’’ method.

A. Using the measurements z

The KLTS method relies on the use of the z_{1k}, z_{2k}, z_{3k} measurements, which are Gaussian random variables (r.v.), instead of the $\hat{\sigma}_1^2, \hat{\sigma}_2^2$ and $\hat{\sigma}_3^2$ elementary estimates, which are a linear combination of χ^2 random variable. The main advantage of this approach lays in the property of the Gaussian r.v. which remain Gaussian when they are linearly combined.

However, these measurements are strongly correlated for two reasons:

- 1) the z_{ABk} and z_{ABk+1} are not independent (except in the case of White FM and AVAR without overlapping);
- 2) for a given k , the z_{1k}, z_{2k} and z_{3k} are not independent (e.g. their sum is null if the measurement noise is neglected).

We will postpone the treatment of the correlation between successive measurements. We first treat the correlations between the 3 measurements at a given time. Hence, let us assume that the $3M$ measurements form M independent triplets, successive realizations of 3 correlated Gaussian r.v. We aim to determine three linear combinations of these r.v. that are independent one of each other:

$$w_{l,k} = \alpha_l z_{1,k} + \beta_l z_{2,k} + \gamma_l z_{3,k}, \quad l = 1, 2, 3 \quad (11)$$

The solution of this problem is given by the Karhunen-Loève (K.L.) transform: the 9 coefficients $\alpha_l, \beta_l, \gamma_l$ form the eigenvectors of the rotation matrix which diagonalizes the covariance matrix of the measurements, whose diagonal elements are given by Eq.(4) and non diagonal elements by Eq.(5). Note that this matrix is singular in the absence of measurement noise, because $z_{BCk} = z_{ACk} - z_{ABk}$: if the

measurement noise is negligible, a 2 by 2 matrix must be used instead of a 3 by 3 matrix.

Now, since the $w_{l,k}$ are all independent, the PDF is easy to calculate in the model world, by assuming definite values for the true variances: $p(\vec{W}|\vec{\Theta}) = \prod_{k=1}^M \prod_{l=1}^3 p(w_{l,k}|\vec{\Theta})$. To render more explicit this expression, let us introduce the three K.L. variances V_l obtained by diagonalization of the covariance matrix. All PDF are Gaussian and their product can be written as:

$$p(\vec{W}|\vec{\Theta}) = \prod_{l=1}^3 \frac{1}{V_l^{M/2}} \exp\left(-\frac{\sum_{k=1}^M w_{l,k}^2}{2V_l}\right) \quad (12)$$

B. Using the sufficient statistic properties of the estimates

In this equation, it is quite interesting to develop, using Eq. (11), the numerator of the exponential argument, which is the only term that depends on the actual measurements:

$$\begin{aligned} \sum_{k=1}^M w_{l,k}^2 &= M \left[\alpha_l^2 \hat{\sigma}_1^2 + \beta_l^2 \hat{\sigma}_2^2 + \gamma_l^2 \hat{\sigma}_3^2 + 2\alpha\beta \widehat{\text{GCov}}_A \right. \\ &\quad \left. - 2\alpha\gamma \widehat{\text{GCov}}_B + 2\beta\gamma \widehat{\text{GCov}}_C \right]. \end{aligned} \quad (13)$$

Eq (13) means that the only knowledge of the 6 elementary or final estimates is sufficient to compute the PDF of the actual set of $3M$ measurements. The number of estimates reduces to 2 in the absence of measurement noise (in this case, the final estimates can be computed from the elementary ones).

This result was expected: indeed, the variance estimates form a sufficient statistics for the variance estimation, meaning that the precise knowledge of the set of measurements $\vec{X} = (x_1, \dots, x_n)^T$ does not bring any new information beyond the estimates. More precisely, the vector of estimates \vec{E} forms a sufficient statistics for $\vec{\Theta}$ if $p(\vec{X}|\vec{\Theta}) = f(\vec{E}, \vec{\Theta}) \cdot g(\vec{X})$, where f and g are two functions [8]. Indeed, we obtain after Bayesian inversion: $p(\vec{\Theta}|\vec{X}) \propto \pi(\vec{\Theta}) \cdot p(\vec{X}|\vec{\Theta}) \propto \pi(\vec{\Theta}) \cdot f(\vec{E}, \vec{\Theta}) \cdot g(\vec{X})$. In the experimenter world, \vec{X} is formed by actual measurements, with known values. Hence, $g(\vec{X})$ appears as a constant. On the other hand, $p(\vec{\Theta}|\vec{X})$ is a function of the random variable $\vec{\Theta}$, whose integral is unity. Because of this normalisation, the multiplication by the constant $g(\vec{X})$ does not change the value of the function, that is entirely determined by $f(\vec{E}, \vec{\Theta})$ and the prior. However, it does not mean that only the PDF of the estimates makes sense. We prove in this paper that the best way is to use these estimates to compute the Gaussian PDF of the data themselves, that remains Gaussian after applying the K.L. transform.

Because we only use the estimates, the correlation between successive data does not lead to more complexity. Indeed, let N be the number of degrees of freedom corresponding to a set of M measurements, with $N < M$. The estimation of N has been reported in [9]. We have simply to write Eq. (12) for N independent measurements by using estimates computed on M correlated measurements, giving:

$$p(\vec{W}|\vec{\Theta}) = \prod_{l=1}^3 \frac{1}{V_l^{N/2}} \exp\left(-\frac{N \sum_{k=1}^M w_{l,k}^2}{2MV_l}\right) \quad (14)$$

C. The KLTS algorithm

To compute $p(\bar{\Theta}|\bar{W})$, we apply the same approach as in [5]. To take into account all the a priori values of $\bar{\Theta}$, we use a Monte-Carlo scheme with random sampling. This sampling ensures the observance of the total ignorance a priori law: the samples are chosen at random on a logarithmic scale, independently for each variance. We work in the experimenter point of view: we assume that two triplets of estimates, final and elementary, have been calculated from the $3M$ elementary measurements. These six numbers have six definite values that will be used in the calculations detailed below. The different steps of the calculation are performed in the same order as in [5] (the common steps of the two methods are recalled here, for sake of completeness):

- Choose at random a triplet of true variances, with a uniform probability on a logarithmic scale for each variance and independence between the three variances.
- Calculate for this triplet the covariance matrix of the measurements given by Eq. (1) and (3).
- Calculate the eigenvectors and eigenvalues of this covariance matrix
- Use these values to calculate the PDF given in Eq. (14), using Eq. (13).
- For an exhaustive exploration of the PDF, repeat 10^7 times the entire process.
- For each of the three variables, normalize the probability densities by dividing by their sum (sum of 10^7 values).
- Also for each of the three variables, sort the true variance values and calculate the cumulative density function by a partial sum on the associated normalized probability densities.
- determine a confidence interval at 95% on each true variance from the corresponding cumulative density functions.
- Verify that the low limit of the confidence interval is meaningful. For a Gaussian distribution, 99.7% of data are included in a confidence interval at $\pm 3\sigma$. If the low limit of this $\pm 3\sigma$ confidence interval (in logarithmic scale) is smaller than the low limit of the a priori range (here 10^{-5}), we suspect (and have verified) that the low limit of the $\pm 2\sigma$ confidence interval calculated in the preceding step will depend on the low limit of the a priori range. If it occurs, we replace the low limit of the confidence interval by 0.

Finally, we calculate the median value, i.e. the argument giving the CDF equal to 0.5. This value, always positive, may be an alternative estimate of the parameters.

IV. VALIDATION OF THE KLTS METHOD BY MONTE-CARLO SIMULATIONS

A. Principle of the simulation

In order to validate the KLTS method, we have compared its results to Monte-Carlo simulations.

The algorithm is as follows:

- 1) Select a target set of 3 final estimates $(\hat{\sigma}_A^2, \hat{\sigma}_B^2, \hat{\sigma}_C^2) = (A_0, B_0, C_0)$. We call it “reference estimate set”.

- 2) Draw at random a parameter triplets $(\sigma_A^2, \sigma_B^2, \sigma_C^2)$.
- 3) Randomly draw $3N$ measurements $\{z_{ABk}, z_{BCk}, z_{CAk}\}$ according to $(\sigma_A^2, \sigma_B^2, \sigma_C^2)$.
- 4) Compute the final estimates $(\widehat{\text{GCov}}_A, \widehat{\text{GCov}}_B, \widehat{\text{GCov}}_C)$ from these measurements.
- 5) If this final estimate set is close to the reference estimate set within 10 %¹, the corresponding parameter triplet $(\sigma_A^2, \sigma_B^2, \sigma_C^2)$ is kept, otherwise it is thrown.
- 6) Go to Step 2).

Each simulation run stops when 10,000 achievements have been obtained in order to provide a meaningful knowledge of the parameter statistical distributions.

This ensemble of 10,000 parameter triplets giving $(\hat{\sigma}_A^2, \hat{\sigma}_B^2, \hat{\sigma}_C^2) = (A_0, B_0, C_0)$ is then compared to the confidence interval obtained by the KLTS method.

Thanks to the sufficient statistic properties of the estimates, it turns out that any $3N$ measurement set providing the reference estimate set leads to the same statistical distribution of the parameter triplet.

B. Results and discussion

We will focus this study only on the efficiency of the method, i.e. its ability to fit the true confidence intervals, and not on the behavior of the confidence intervals versus different parameters. More details on this latter subject may be found in [5].

This validation way checks the accuracy of the KLTS method with respect to several physical variables.

1) *Influence of the EDF*: In order to compare the confidence interval given by several other methods, we chose to set the final estimate values to $\hat{\sigma}_A^2 = \hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$ and to vary the number of EDF. The results given by the different methods, a simple Gaussian approximation [4], the Ekstrom-Koppang method (EK) [3] and the KLTG method [5] are plotted in Figure 2-A. All these confidence intervals fit pretty well the empirical error bars obtained by 10,000 Monte-Carlo simulations above 200 EDF. Between 50 and 200 EDF, EK and KLTG remain usable but only KLTG fits quite well between 5 and 50 EDF. Finally, below 5 EDF, none of these methods is satisfactory.

The new KLTS method is compared to KLTG in Figure 2-B. The KLTG confidence intervals fit pretty well the empirical error bars even for 2 EDF. The same is true for the median estimates which correspond perfectly to the empirical median values at the center of the error bars.

However, this comparison is limited to 2 EDF because the set $\hat{\sigma}_A^2 = \hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$ is impossible for 1 EDF without measurement noise. It was demonstrated that, if the measurement noise is negligible, and this is always the case when there is only 1 EDF, the 3rd final estimate is totally determined by the other 2 ones [5]. For this reason, we set $\hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$ which leads to $\hat{\sigma}_A^2 = -0.5$. Table I compares

¹This 10 % interval was selected by a compromise ensuring that the final estimate set is sufficiently close on a log scale to the reference estimate set, while the calculation time is reasonable (the computation time increases as the cube of the inverse of the interval width).

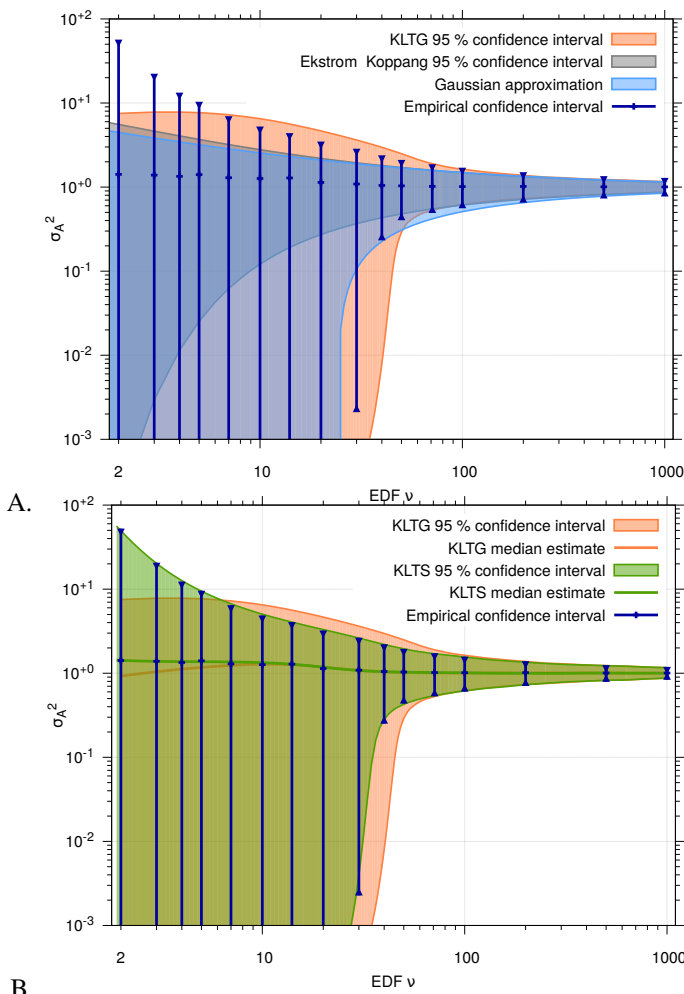


Fig. 2. Estimation of the 95 % confidence interval for a set of clocks with equal final estimates ($\hat{\sigma}_A^2 = \hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$) versus the EDF number. The graph above (A) shows the results obtained from previous methods and the graph below (B) presents the results of the KLTS method (green area) compared to the KLTG (red area). The reference is given by the blue error bars which were obtained by massive Monte-Carlo simulations (see algorithm §IV-A).

TABLE I

COMPARISON FOR 1 EDF OF THE CONFIDENCE INTERVALS AS WELL AS THE MEDIAN ESTIMATES (50 %) OBTAINED BY THE MONTE-CARLO SIMULATIONS (EMPIRICAL), BY THE KLTG METHOD OF [5] AND BY THE NEW KLTS METHOD. THE (A_0, B_0, C_0) ESTIMATE TRIPLET IS $(-0.5, 1, 1)$.

Bounds	$\hat{\sigma}_A^2 = -0.5$		$\hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$	
	Emp.	2.5 %	$1.78 \cdot 10^{-5}$ (2.5 %)	$3.0 \cdot 10^{-5}$ (2.5 %)
	50 %	0.192 (50.0 %)	76 (95.0 %)	182 (97.5 %)
	95 %	27.5 (95.0 %)	1.23 $\cdot 10^{-4}$ (5.31 %)	0.42 (37.0 %)
	97.5 %	92 (97.5 %)	3.23 (70.4 %)	5.3 (76.3 %)
KLTG	2.5 %	$1.36 \cdot 10^{-5}$ (1.52 %)	0.90 (50.0 %)	3.23 (70.4 %)
	50 %	$5.5 \cdot 10^{-3}$ (28.7 %)	76 (95.0 %)	5.3 (76.3 %)
	95 %	0.83 (68.8 %)	1.23 $\cdot 10^{-4}$ (5.31 %)	0.42 (37.0 %)
	97.5 %	1.38 (75.0 %)	3.23 (70.4 %)	5.3 (76.3 %)
KLTS	2.5 %	$1.67 \cdot 10^{-5}$ (2.59 %)	0.90 (50.0 %)	3.23 (70.4 %)
	50 %	0.200 (50.6 %)	76 (95.0 %)	182 (97.5 %)
	95 %	35 (95.5 %)	0.90 (49.8 %)	90 (95.3 %)
	97.5 %	98 (97.8 %)	208 (97.6 %)	




Fig. 3. Estimation of the 95 % confidence intervals for 2 clocks with equal final estimates ($\hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$) versus the final estimate of the other clock ($-0.5 \leq \hat{\sigma}_A^2 \leq 100$). The number of EDF is 2. The graph above (A) shows the results obtained from positive final estimates on a log-log plot whereas the graph below (B) uses a linear scale in order to display also the results obtained from negative final estimates. The error bars, red for σ_A^2 and green for σ_B^2 and σ_C^2 , were obtained by massive Monte-Carlo simulations.

the confidence interval obtained by KLTG and KLTS to the empirical bounds given by 10,000 Monte-Carlo simulations.

Unlike the results of KLTG, the KLTS bounds as well as the median estimate show a good agreement with the empirical bounds and median. However, because of the very low level of the 2.5 % bounds, they should be considered as equal to 0. In these conditions, KLTS gives fully reliable 95 % upper limits (95.7 % for σ_A^2 , 95.6 % for σ_B^2 and σ_C^2).

2) *Influence of the disparity of the 3 final estimates:* Figure 3 displays, for 2 EDF, the evolution of the confidence interval of all parameters versus the final estimates $\hat{\sigma}_A^2$ which varies from 0.01 to 100 whereas the 2 other ones are set to $\hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$. On the log-log plot as well as on the linear plot, the agreement between the error bars and the KLTS confidence intervals is excellent.

3) *Influence of the measurement noise:* The main advantage of GCov over the three-cornered hat relies in its rejection of the measurement noise. Since the measurement noise is directly addressed by KLTS, e.g. to solve equation (13), it




Fig. 4. Estimation of the 95 % confidence intervals for a set of clocks with equal noise measurements ($\hat{\sigma}_A^2 = \hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$) versus the measurement noise level ($0.01 \leq \sigma_N^2 \leq 50$). The number of EDF is 100. The results obtained by the KLTS method (green area) may be compared to the ones obtained by KLTG (red area). The blue error bars were obtained by massive Monte-Carlo simulations.

is of importance to study its influence on the confidence intervals. This influence is shown on Figure 4 where the final estimates are set to ($\hat{\sigma}_A^2 = \hat{\sigma}_B^2 = \hat{\sigma}_C^2 = 1$), the number of EDF is 100 and the variance of the measurement noise σ_N^2 varies from 0.01 to 50. Despite slight discrepancies between the confidence intervals and the error bars obtained from the Monte-Carlo simulations, the agreement is quite good. The main differences appear for $1 < \sigma_N^2 < 10$ when the 2.5 % bound decreases drastically. This effect is also visible for the median estimate which seems to decrease at a different rate for the theoretical and the empirical median estimates. A very small discrepancy is also visible for the upper bound at $\sigma_N^2 = 7, 10$ and 20 . These discrepancies can be explained as follows: in the PDF computation of KLTS, the product (14) of the elementary probabilities of each data becomes extremely small when the number of data becomes large, and the result of (14) can be wrongly truncated to 0 even in double precision. As a consequence, the use of KLTS must be restricted for cases where the number of EDFs does not exceed one hundred. For hundred EDFs as in Figure 4, the truncation errors already explain some discrepancies for a high level of measurement noise, that do not however affect much the upper bound of the confidence interval, undoubtedly the most important result.

4) *Discussion: KLTS vs KLTG:* KLTS is undoubtedly a rigorous approach which does not rely on approximations or simplifying assumptions. As a consequence, it provides very relevant confidence intervals even for very low EDF, including the limit case of 1 EDF (see Table I), as proved by the almost perfect agreement between the confidence intervals given by KLTS and obtained by the Monte Carlo simulations in Figures 2 to 4. The only slight differences can be attributed to the way of computing the PDF by numerical integration.

However, we have seen that truncation errors can occur with KLTS when the number of data becomes large: the use of KLTS must be restricted for cases where the number of EDFs

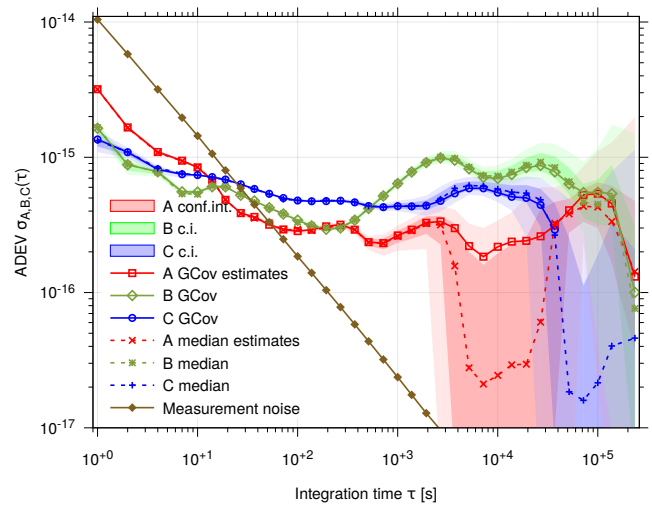


Fig. 5. GCov measurements of a set of 3 Cryogenic Sapphire Oscillators of FEMTO-ST. The confidence intervals are represented by colored areas: bright colors for 68 % c.i. (1σ) and pale colors for 95 % c.i. (2σ). The brown line shows the measurement noise obtained from the closure. The linear frequency drifts of the CSO have been removed.

does not exceed one hundred.

On the other hand, KLTG relies on a Gaussian approximation of the estimates [5] which can be assumed only for high EDF. Unlike KLTS, KLTG must be restricted for high EDFs. Figure 4 shows that KLTG and KLTS provide almost the same results for 100 EDFs and these results are in perfect agreement with the Monte-Carlo simulations. KLTG is then a very good substitute to KLTS above 100 EDFs.

C. Application to a set of real clocks

These methods were applied to assess the confidence intervals for the GCov measurements of a set of 3 Cryogenic Sapphire Oscillators designed and made in FEMTO-ST. These measurements were carried out by the Tracking DDS designed and made in INRIM (for more details, see [10]). It may be noticed that, since the clocks are in the same room, we can not fully assume their total independence. Figure 5 presents the same measurements as the ones plotted in Figure 11 (below) of [10] but we have added the 95 % confidence intervals obtained by using KLTS (below 100 EDF, i.e. $\tau > 5000$ s) and KLTG (above 100 EDF, i.e. $\tau < 5000$ s).

As expected, the confidence intervals are pretty tight around the estimates for low τ values whereas they extend downward for high τ values. The lower bounds tend toward 0 above $\tau = 2000$ s for CSO A, above $\tau = 20000$ s for CSO C and above $\tau = 100000$ s for CSO B. In these cases, only an upper limit of the stability of the clock may be assessed. As expected, the less stable clocks have the smallest confidence intervals and we can see these intervals increasing or decreasing regarding the relative positions of the other CSOs. For instance, A is the most stable clock between 500 s and 30000 s but becomes the less stable clock around 10^5 s, as a consequence, the lower bound of its confidence interval increases from ~ 0 at 10^4 s to $\sim 2 \cdot 10^{-16}$ at 10^5 s.

The median estimates remain generally close to the GCov estimates. However, when the lower bound of a clock tends

toward 0, the median estimate decreases significantly and becomes far lower than the GCov estimate when this latter exists, i. e. when it is positive. Nevertheless, there is still a positive median estimate even when the GCov estimate is negative (e.g. C above $\tau = 30\,000$ s). In such a case the upper limit of the confidence interval is the relevant information, while the estimate is of little value. It may be noticed that the relevance of the median estimates could be improved by using a more stringent prior $\pi(\vec{\Theta})$ in Eq. (10), i.e. a prior based on a perfectly objective *a priori* knowledge of the range in which the parameters can vary. To stay in the most general case, we have preferred to stick to a “total ignorance” prior (see III-C) in this paper.

We have also added the ADEV of the measurement noise on Figure 5. It is prevailing below 10 s and we can see that the confidence interval of the most stable clock in this range, B and C, have a larger confidence interval (e.g. C at 1 s) than for $\tau = 100$ s for instance. At 1 s, the measurement noise is approximately 5 times higher than the stability of C and, despite its huge EDF number (EDF = 370 000 @ 1 s), the effect of the measurement noise affects clearly its confidence interval as in Figure 4. In other words, the measurement noise is too high to be fully rejected by GCov.

V. CONCLUSION

The method we propose, KLTS, provides the Cumulative Density Functions of the clock noise variances and the corresponding confidence intervals. Unlike the direct GCov estimates, the medians of these intervals cannot be negative.

KLTS is a rigorous method since it involves neither approximations nor simplifying assumptions. It is based on the property of sufficient statistics that form the estimates.

KLTS is valid for very low EDF including 1 EDF. Massive Monte-Carlo simulations have perfectly validated KLTS in different contexts: confidence interval versus the Equivalent Degrees of Freedom, vs the the stability disparity of the clocks and vs the level of the measurement noise.

However, KLTS suffers from the drawback of not being easily computed for high EDF (> 100). But it has been proved in this paper that the KLTG method, presented in a previous paper [5], is perfectly reliable above 100 EDF even if the measurement noise is strong. The combination of these two methods provides then a powerful tool to assess confidence intervals for the clock noise variances, whatever the EDF and the measurement noise level.

ACKNOWLEDGEMENT

This work was partially funded by the ANR Programmes d’Investissement d’Avenir (PIA) Oscillator IMP (Project 11-EQPX-0033) and FIRST-TF (Project 10-LABX-0048) as well as by the Conseil Régional de Bourgogne Franche-Comté.

REFERENCES

[1] J. Gray and D. Allan, “A method for estimating the frequency stability of an individual oscillator,” in *Proceedings of the 28th Annual Symposium on Frequency Control*, May 1974, pp. 243–246.

[2] J. Gros Lambert, D. Fest, M. Olivier, and J. Gagnepain, “Characterization of frequency fluctuations by crosscorrelations and by using three or more oscillators,” in *Proceedings of the 35th Frequency Control Symposium*, May 1981, pp. 458–463.

[3] C. R. Ekstrom and P. A. Koppang, “Error bars for three-cornered hats,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 53, no. 5, pp. 876–879, 2006.

[4] F. Vernotte, C. E. Calosso, and E. Rubiola, “Three-Cornered Hat versus Allan Covariance,” in *IEEE International Frequency Control Symposium (IFCS)*, May 2016, pp. 1–6.

[5] F. Vernotte and E. Lantz, “Three-cornered hat and Gros Lambert covariance: A first attempt to assess the uncertainty domains,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 3, pp. 643–653, March 2019.

[6] A. Premoli and P. Tavella, “A revisited three-cornered hat method for estimating frequency standard instability,” *IEEE Transactions on Instrumentation and Measurement*, vol. 42, no. 1, pp. 7–13, February 1993.

[7] A. S. Kholevo, *Encyclopedia of Mathematics*. Springer Science+Business Media B.V. / Kluwer Academic Publishers, 2001, ch. Sufficient statistic, ISBN 978-1-55608-010-4. [Online]. Available: https://www.encyclopediaofmath.org/index.php/Sufficient_statistic

[8] G. Saporta, *Probabilités Analyse des Données et Statistiques*. Paris: Editions Technip, 1990.

[9] C. Greenhall and W. Riley, “Uncertainty of stability variances based on finite differences,” in *35th annual Precise Time and Time Interval meeting*, San Diego (California, USA), December 2003, pp. 267–280.

[10] C. E. Calosso, F. Vernotte, V. Giordano, C. Fluhr, B. Dubois, and E. Rubiola, “Frequency stability measurement of cryogenic sapphire oscillators with a multichannel tracking DDS and the two-sample covariance,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 3, pp. 616–623, March 2019.

APPENDIX: GLOSSARY OF SYMBOLS

A. Measurements

\bar{y}_{ABk}	k^{th} frequency deviation sample between clock A and clock B
\bar{y}_{1k}	k^{th} frequency deviation sample at the output of counter 1
\bar{y}_{N1k}	k^{th} frequency deviation sample of the intrinsic measurement noise of counter 1
z_{ABk}	Difference between the two consecutive clock frequency deviation samples \bar{y}_{ABk+1} and \bar{y}_{ABk} divided by $\sqrt{2}$
z_{1k}	Difference between the two consecutive counter frequency deviation samples \bar{y}_{1k+1} and \bar{y}_{1k} divided by $\sqrt{2}$
z_{N1k}	Difference between the two consecutive noise frequency deviation samples \bar{y}_{N1k+1} and \bar{y}_{N1k} divided by $\sqrt{2}$
$z_{cls,k}$	Sum of the k^{th} z measurements of all counters (closure relationship)

B. Parameters and estimates

The list below gives the parameters, i.e. the mathematical expectation we want to assess. The same symbols with a hat $\hat{\cdot}$ stand for the estimates of these parameters, i.e. the values computed from a finite number of measurements.

σ_{AB}^2	Allan variance of the comparison between clock A and clock B
σ_1^2	Allan variance of the output of counter 1 (elementary estimate)
σ_{N1}^2	Allan variance of the intrinsic measurement noise of counter 1. Since the noise level is assumed to be the same for all counters, the subscript 1, 2 or 3 may be omitted.
σ_{cls}^2	Allan variance of the closure
σ_A^2	Allan variance of clock A (σ_A^2 is called model parameter whereas $\hat{\sigma}_A^2$ is called final estimate)
TCH $_A$	Three-cornered hat computation of Allan variance for clock A
GCov $_A$	Gros Lambert covariance computation of Allan variance for clock A

C. Bayesian formalism

$p(\vec{X} \vec{\Theta})$	Conditional probability density function of the multidimensional measurement \vec{X} knowing the multidimensional parameter $\vec{\Theta}$ (model world)
$p(\vec{\Theta} \vec{X})$	Conditional probability density function of the multidimensional parameter $\vec{\Theta}$ knowing the multidimensional measurement \vec{X} (experimenter world)
$\pi(\vec{\Theta})$	<i>A priori</i> probability density function of the multidimensional parameter $\vec{\Theta}$ before any measurement (prior)

D. KLTS method

$\begin{pmatrix} \alpha_1 & \beta_1 & \gamma_1 \\ \alpha_2 & \beta_2 & \gamma_2 \\ \alpha_3 & \beta_3 & \gamma_3 \end{pmatrix}$	Eigenvectors of the rotation matrix which diagonalizes the covariance matrix of the measurements (Karhunen-Loève Transform)
$\vec{\Theta} = (V_1, V_2, V_3)^T$	KLT variances, i.e. eigenvalues of the covariance matrix of the measurements
$\vec{W} = (w_{1k}, w_{2k}, w_{3k})^T$	Image of the measurements $(z_{1k}, z_{2k}, z_{3k})^T$ by the KLT
M	Number of measurements
N	Number of Equivalent degrees of freedom