https://doi.org/10.23939/sisn2019.01.053

UDC 004.4, 519.688

ANALYSIS OF THE INFLUENCE OF SAMPLE RATES ON THE ALLAN VARIANCE

T. Marusenkova,
Lviv Polytechnic National University, 12, S. Bandery Str., Lviv, Ukraine,
E-mail: tetyana.marus@gmail.com, ORCID: 0000-0003-4508-5725

The paper considers the problem of interpreting the Allan deviation plot for signals from sensors polled more frequently than data are refreshed. The Allan variance is a standard tool for analysis of noise terms inevitably present in signals of inertial sensors. There exists a well-defined algorithm for its calculation both for time domain and frequency domain. Having calculated the Allan variance as a function of time (or frequency) one fetches its square root, called Allan deviation, and builds its plot in a logarithmic format. Each region of the Allan deviation plot characterizes a specific noise kind (white noise, flicker noise, random walk, etc). The plot is expected to have a well recognizable, predefined shape. However, in practice it may be that a plot obtained for real time series does not follow its textbook pattern. In this case it is unobvious how to interpret the plot and whether it is applicable or not. We observed quite untypical Allan deviation plots for signals of a magnetometer sampled too frequently, which suggested that the sample rate can be responsible for the unusual shape of the plot. Our work is aimed at analyzing the influence of the sample rate on the Allan deviation plot and evaluating the applicability of such a plot obtained for signals sampled too frequently. We reproduced experimental results by simulation and detected that the sample rate for synthesized white noise signals impacts the shape of the Allan deviation plot. The same idea was corroborated by filtering out repeated measurement points from experimentally obtained magnetometer signals. The simulation results are backed up by analytical calculations. Therefore, all the applied approaches such as simulation, filtering reading of a real sensor and analytical considerations confirmed that the shape of the Allan deviation plot depends on the signal sample rate. Moreover, we have shown that the Allan deviation plot built under these conditions is completely inapplicable unless all repeated measurement points are filtered out. Our analytical explanation of this fact is confirmed by a set of experiments. We provide a detailed description of a procedure for evaluation of the applicability of the Allan deviation plot using a magnetometer.

Key words: Allan variance, noise, inertial sensor, sample frequency, magnetometer.

АНАЛІЗ ВПЛИВУ ЧАСТОТИ ДИСКРЕТИЗАЦІЇ НА ДИСПЕРСІЮ АЛАНА

Т. Марусенкова,
НУ “Львівська політехніка”, кафедра програмного забезпечення, Україна,
м. Львів, вул. С. Бандери, 12, E-mail: tetyana.marus@gmail.com, ORCID: 0000-0003-4508-5725

Розглянуто проблематику інтерпретації графіка девіації Алана для сигналів сенсорів за умови, що частота дискретизації перевищує швидкість оновлення показів. Варіація Алана є стандартним інструментом аналізу шумових складових, неминуче присутніх у сигналах будь-яких інерційних сенсорів. Існує повністю визначений алгоритм розрахунку варіації Алана як для часової, так і для частотної областей. Після визначення варіації Алана як функції часу (або частоти) розраховують девіацію Алана (квадратний корінь варіації Алана) і будують її графік у логарифмічному форматі. Кожна ділянка цього графіка характеризує шум певного типу (білий шум, рожевий

шум, випадкове блукання тощо). Очікується, що форма графіка девіації Алана загалом відповідає встановленому взірцю і може бути легко розпізнана. Однак на практиці форма графіка може істотно відрізнятися від книжкового шаблону. У такому випадку стає неочевидним, як інтерпретувати графік і чи він взагалі є придатним до використання. Ми спостерігали нетипові графіки девіації Алана для сигналів магнітометра, отриманих з частотою дискретизації, що перевищувала швидкість оновлення показів, завдяки чому виникла ідея про залежність форми графіка від частоти дискретизації. Метою статті є аналіз впливу частоти дискретизації на форму графіка девіації Алана та оцінювання придатності цього графіка за неправильно вибраної частоти дискретизації. Проведене нами імітаційне моделювання дало змогу якісно відтворити експериментальні результати. Показано, що частота дискретизації згенерованого білого шуму впливає на форму графіка. Цей самий висновок дозволяє зробити і фільтрування повторених точок вимірювання з реальних сигналів магнітометра. Наведено аналітичні розрахунки, що пояснюють і підтверджують вплив частоти дискретизації на форму графіка. Нами показано, що графік девіації Алана для сигналів, отриманих з частотою дискретизації, що перевищує швидкість оновлення показів, не придатний для застосування, якщо не фільтрували повторені точки вимірювання. Аналітичне пояснення цього факту підтверджено експериментально. Подано детальний опис процедури оцінювання придатності графіка девіації Алана за допомогою магнітометра.

Ключові слова: дисперсія Алана, шум, інерційний сенсор, частота дискретизації, магнітометр.

Problem statement

Inertial measurement units (IMUs) especially those based on micro-electromechanical systems (MEMS) find their application in various spheres including but not limited to medicine [1, 2], avionics and robotics. For instance, inertial sensors are used in order to recognize a human posture/gait, distinguish between normal and pathological human movements, classify the actions of a worker as right or wrong, track the movements of a sportsman, and even detect the first signs of loosing footage and prevent damages caused by falling down [3]. Their key feature is that they are autonomous, i.e., they do not require any external information for estimation of the coordinates and orientation, in contrast to well-known GPS navigation systems. However, MEMS IMUs are susceptible to errors, which should be carefully modeled, studied and compensated for, otherwise these devices are scarcely applicable.

Typically an IMU contains a triaxial accelerometer, a triaxial gyroscope and additionally a 3D magnetometer. IMU devices supplied with a magnetometer are commonly called MARG (Magnetic, Angular Rate and Gravity). In order to achieve more consistent and reliable results, they use data fusion [4, 5, 6] – readings of all the sensors are combined to fight fallacies of each individual sensor. A gyroscope measures angular velocity that, when integrated over time, produces the sensor orientation (in the sensor frame, not the earth frame). MEMS gyroscopes are prone to errors, which quickly cumulate due to integration. Consequently, a gyroscope alone is not able to provide an absolute measurement of orientation. An accelerometer measures the earth's gravitational field mingled with accelerations due to motion. For this reason, an accelerometer cannot be used alone to provide the absolute attitude as well. Finally, a magnetometer will measure the earth's magnetic field mixed with a local magnetic field and possible distortions, which means that it cannot be used as the only tool for measuring the absolute orientation, too. Instead, data obtained from these three devices are used together. Each of these devices is susceptible to a number of errors, which fall into two groups: deterministic and stochastic.

A standard tool of characterizing noise terms (stochastic errors) in signals of inertial sensors is the Allan variance or, to be more precise, the Allan deviation plot. Each part of this plot is responsible for a specific kind of noise (for example, quantization noise, white noise, bias, rate ramp, etc). One expects that the Allan deviation plot for a time history collected from a real inertial sensor resembles the academic Allan deviation plot commonly shown in the subject literature. In practice, however, the Allan deviation

plot may be hardly recognizable, and in this case it becomes unobvious how to interpret the plot and whether it is applicable or not.

The work deals with the case of a magnetometer polled more frequently than it actually could provide new readings and shows the relationship between the sample frequency and the Allan deviation plot applicability.

Recent research and publications analysis

The Allan variance method was originally developed by David W. Allan for analysis of frequency stability in oscillators (clocks), but now it is widely used for processing other data, including readings of inertial sensors. There are several modifications of this method, the overlapping Allan variance method [7, 8] being the most general and universal. It is composed of the following steps.

1. Keeping a sensor steady, one acquires an equidistant time series, *y(t)*, of length *N*, with the sample period τ_0. If the collected time series turns out to be non-equidistant, its analysis would be much more complicated [9].

2. For some averaging factor *m*, one takes all possible overlapping sample clusters of period $\tau = m\tau_0$. For example, for $m = 3$, the first sample cluster contains readings #1, #2, #3 and #4, with the total time elapsed between readings #1 and #4, equal to $3\tau_0$. The second sample cluster is comprised by readings #2, #3, #4 and #5, and so on. The last cluster for the chosen *m* contains readings #(*N*–3), #(*N*–2), #(*N*–1) and #*N*. Each pair of neighboring sample clusters is always separated by the sample period, τ_0. The process is depicted in Fig. 1, a.

3. Once all the possible $(N - m)$ sample clusters have been formed for the selected value *m* (and τ), the Allan variance is calculated as a function of τ. There are several well-known formulas, which can be transformed into each other. The most common formulas are:

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2 (N - 2m)} \sum_{i=1}^{N-2m} [x_{i+2m} - 2x_{i+m} + x_i]^2 \quad (1)$$

$$\sigma_y^2(\tau) = \frac{1}{2m^2 (N - 2m)} \sum_{j=1}^{N-2m} \left\{ \sum_{i=j}^{j+m-1} [\bar{y}_{i+m} - \bar{y}_i] \right\}^2 \quad (2)$$

$$\sigma_y^2(\tau) = \frac{1}{2m^2 (N - 2m)} \sum_{i=1}^{N-2m} \left[-\sum_{j=i}^{i+m-1} \bar{y}_j + \sum_{j=i+m}^{i+2m-1} \bar{y}_j \right]^2 \quad (2a)$$

Here x_i $(i = \overline{1, N})$ represent the cumulated sum of the collected time history y_i $(i = \overline{1, N})$. In practice they usually use (1) because (2) consumes many more computational resources due to lots of summation operations. In order to use (1) one should integrate the time history of measurements first. In [10] detailed interim calculations can be found.

4. Calculate the Allan deviation for the value of τ being currently considered. The Allan deviation is the square root of the Allan variance.

5. Repeat steps 2–4 for different values of *m* (and τ) and draw the Allan deviation plot, usually in a log-log format. Different parts of the plot allow us to figure out different characteristics of the noise terms in a signal being analyzed. Few literature sources focus on the choice of the averaging factor, *m*. On the contrary, many of them do not specify how many values of *m* should be taken to ensure meaningful results of the Allan variance analysis. Moreover, there are discrepancies between scarce recommendations available when it comes to choosing the minimum value of *m*. Most resources simply start with $m = 1$, however, in [9, 11] it is stated explicitly that *m* should not be less than 9, otherwise averaging would be pointless. In [9] it's said that also no less than 9 different values of *m* (and, correspondingly, τ) should be taken for an Allan deviation plot. Thus, a time history should contain at least 81 items. Several sources assume that an inertial sensor should be kept in an environment free of temperature changes and disturbances overnight, which implies thousands of measurements. For instance, the manual of

ADXRS450/ADXRS453, a flexible inertial sensor evaluation platform by Analog Devices, recommends taking at least 16 000 data points in order to ensure meaningful analysis results. Values of *m* may be consecutive or not, i.e. one can choose to deal with even or odd values of *m* only or with values that are multiples of some natural number other than 1.

The textbook case of the Allan deviation plot [12] reprinted in many sources is shown in Fig. 1, b.

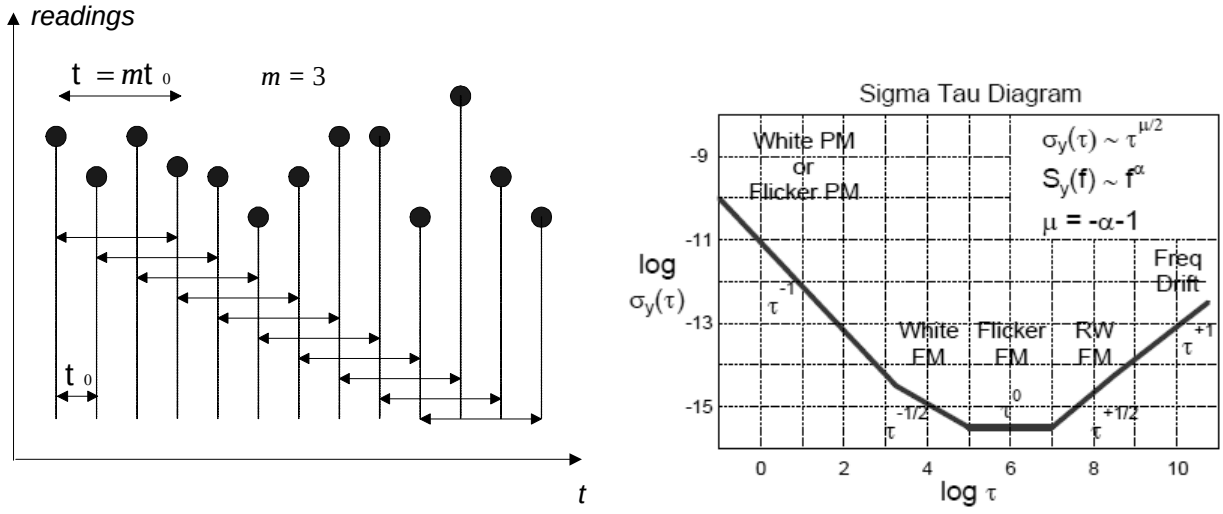

a *b*

Fig. 1. Overlapping sample clusters of sensor readings (a) and the textbook case of the Allan deviation plot (b)

It is common to interpret the Allan deviation plot as follows. The *Y*-value of the leftmost point in the plot represents the noise deviation if no oversampling and no averaging is used. I.e., each measurement point is used alone ($m = 1$) in contrast to the situation when multiple points are taken and averaged together in order to produce a single measurement result. This noise deviation should coincide with the standard deviation calculated in the same way as if the time series in question were a sample, regardless of the order of its items. That is because each point is assessed individually instead of being grouped. As the value of *m* grows, more sensor readings participate in averaging and the averaged value gets closer to the true measurement result, because noise usually has a nearly zero mean when considered over time. This concerns only measurements where high-frequency white noise dominates. Otherwise, the points of the Allan deviation plot show the deviation of something else, different from noise. Thus, the slope in the Allan deviation plot that corresponds to growing *m* values indicates the trade-off between the affordable noise level and time spent on getting a single result by averaging multiple measurement points. I.e., some applications may prefer quick but noisy results, some are targeted at higher accuracy and lower rates. Then the Allan deviation plot reaches its minimum. The *Y*-value of this part of the plot means the least theoretically achievable noise deviation, whereas the corresponding *X*-value suggests how many measurements per second one should take in order to reach this minimal noise level. After reaching its minimum a plot may rise again, due to low-frequency noise such as random walk.

In practice, however, interpretation of the Allan deviation plot may be not so straightforward. The author is a co-developer of IMUTester, a hardware-software tool for exploring IMUs. The tool utilizes a module GY-80 that contains a triaxial digital accelerometer ADXL345, a 3D digital gyroscope L3G4200D, a 3D compass HMC5883L and some auxiliary nodes. Fig. 2 shows data related to the magnetometer and the corresponding Allan deviation plot. The tool is not a focus of this work and is given here only to show that the Allan deviation plot obtained for a magnetometer is similar to neither the textbook case shown in Fig. 1, b nor the plot for white Gaussian noise [13]. When the Allan deviation plot indicates (quite unnaturally) less noise deviation for single measurement points than for groups of averaged points, these issues need additional study and explanation.

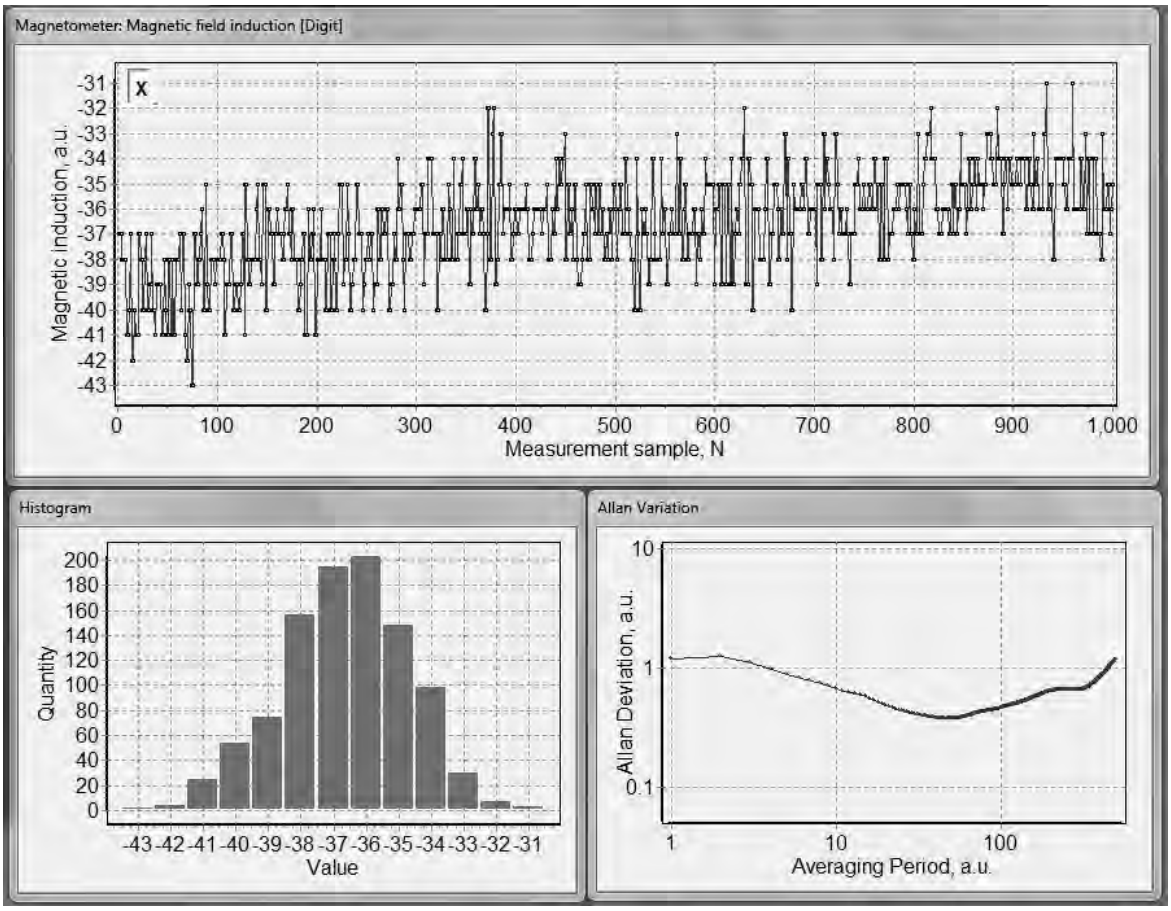

Fig. 2. The Allan deviation obtained for a magnetometer HMC5883L and the magnetic flux density

Statement of purpose

The work is aimed at investigation into the relations between signal sampling rates and the Allan deviation plot. Based on the detected relationships, practical results should be grounded. Moreover, the question about applicability of the corresponding Allan deviation plot should be answered.

A relationship between the sample frequency and Allan deviation

As can be seen on the visualized magnetometer time series, the magnetometer was polled more frequently than it could provide any fresh data, which resulted in repetition of essentially the same measurement results. Our idea is that the shape of the Allan deviation plot is attributable to discrepancies in data refreshment rate and sensor polling rate.

In order to verify quickly whether this idea might have been correct we took the following steps. First we generated a set of time series of length 1000, 3000, 5000, 10000 and 20000 that represented white Gaussian noise of power 1 dB.

Then we replicated each time series so as to obtain a time series where each even item is the copy of its neighbor to the left (i.e., each item of the original was doubled in the new time series). Artificial doubling of measurement points emulates the situation when fresh data are provided twice less frequently than they are sampled ($F_1/F_2 = 2$ where F_1 is the polling rate, F_2 is the data refreshment rate). Then we continued the same process to get time series with each point replicated once more, i.e. $F_1/F_2 = 3$ and so on. We built the Allan deviation plots for each simulated time series.

A sample result is depicted in Fig. 3. The Allan deviation plot for unchanged white noise is in continuity with similar plots for noise of this color that can be found in [13]. One the contrary, the Allan deviation plot for any time series with replicated measurement points was dissimilar to that one for the original simulated white noise but resembled the Allan deviation plot for our magnetometer. Having repeated generation of white noise multiple times for a range of frequencies (and sample periods), we

detected that the results shown in Fig. 3 are qualitatively reproducible. Namely, all plots have a part increasing to some maximum *Y*-value, then decline. Obviously, the more repetitions a time series contains, the more measurement points one needs to take in order to reach the peak *Y*-value.

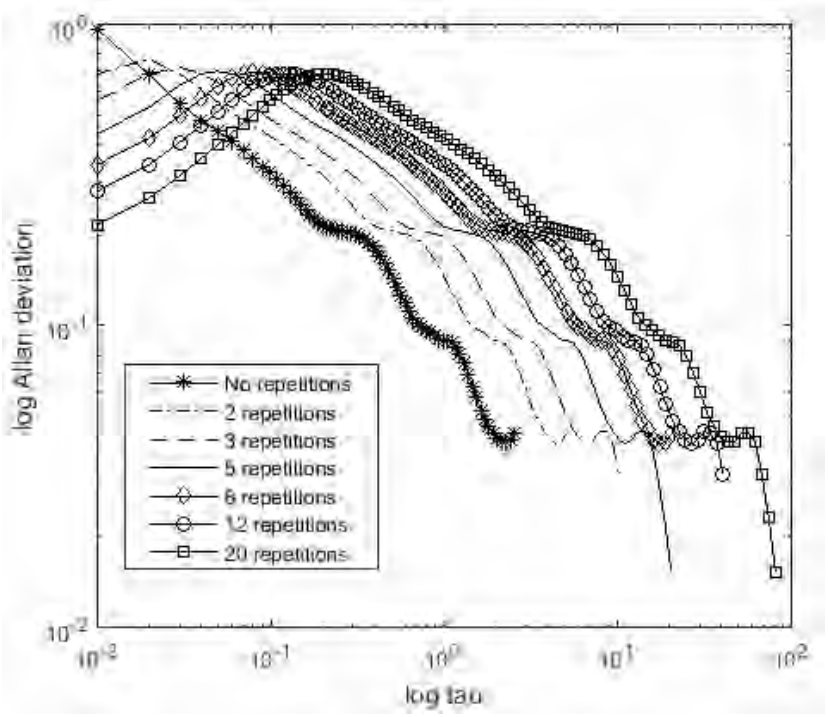

Fig. 3. The Allan deviation plots for simulated white Gaussian noise with repetitions

Consistent results were obtained after filtering the raw magnetometer readings. Namely, we constructed a modified time series from each of original magnetometer signals, having replaced each group of several identical consecutive measurements by a single measurement and reduced the total amount of measurement points, correspondingly. The Allan deviation plot for magnetometer readings filtered in this way resembled the Allan plot for white noise.

Therefore, two opposite approaches, injection of replicated measurement points in simulated white noise signals and removal of repeated points in real magnetometer signals, led to consistent results.

After this preliminary confirmation of the idea about the influence of an improperly chosen sampling rate on the modified shape of the Allan variance, we proceeded with analytical dependencies corroborating this idea. Let us say that *y* is the time history of a sensor whose readings are acquired as frequently as they appear, $y = \{y_1, y_2, y_3, \ldots, y_N\}$. Let us construct several modified signals that contain repetitions. A time series with each item repeated once will be denoted $y2 = \{y_1, y_1, y_2, y_2, y_3, y_3, \ldots y_N, y_N\}$. A time series whose items are replicated twice can be written as $y3 = \{y_1, y_1, y_1, y_2, y_2, y_2, \ldots. y_N, y_N, y_N\}$.

If the amount of repeated consecutive measurement points is termed *R*, then the *i*-th item in the cumulative sum of a time series characterized by *R* can be represented as

$$x_i = \sum_{j=1}^{\hat{e}i/R\acute{u}} Ry_j + (i\%R)\, y_{\hat{e}i/R\acute{u}+1} \tag{3}$$

Thus

$$x_{k+2m} - 2x_{k+m} + x_k = \tag{4}$$

$$= \sum_{j=1}^{\hat{e}(k+2m)/R\acute{u}} Ry_j + ((k+2m)\%R)\, y_{\hat{e}(k+2m)/R\acute{u}+1} - 2\sum_{j=1}^{\hat{e}(k+m)/R\acute{u}} Ry_j - 2((k+m)\%R)\, y_{\hat{e}(k+m)/R\acute{u}+1} + \sum_{j=1}^{\hat{e}k/R\acute{u}} Ry_j - ((k\%R)\, y_{\hat{e}k/R\acute{u}+1}$$

Rearranging sums, one can transform (4) into:

$$x_{k+2m} - 2x_{k+m} + x_k = \tag{5}$$

$$= -\sum_{j=\lfloor k/R \rfloor+1}^{\lfloor (k+m)/R \rfloor} Ry_j + \sum_{j=\lfloor (k+m)/R \rfloor+1}^{\lfloor (k+2m)/R \rfloor} Ry_j + \big((k+2m)\%R\big) y_{\lfloor (k+2m)/R \rfloor+1} - 2\big((k+m)\%R\big) y_{\lfloor (k+m)/R \rfloor+1} + \big(k\%R\big) y_{\lfloor k/R \rfloor+1}$$

It is easy to see that when $R = 1$ (5) reduces to $\left(-\sum_{j=k+1}^{k+m} y_j + \sum_{j=k+m+1}^{k+2m} y_j\right)$ in which one can recognize the radicand of the conventional formula (2a). Let us perform the following transformations:

$$\left\lfloor \frac{k+m}{R} \right\rfloor = \frac{k}{R} + \frac{m}{R} - \frac{(k+m)\%R}{R}, \left\lfloor \frac{k}{R} \right\rfloor = \frac{k}{R} - \frac{k\%R}{R}, \left\lfloor \frac{k+2m}{R} \right\rfloor = \frac{k}{R} + \frac{m}{R} + \frac{m}{R} - \frac{(k+2m)\%R}{R} \tag{6}$$

Since $(x\%R)/R < 1$ for any x, the first two terms in (5) would be zero for any $m < R$, which explains why for m=1 the greater R, the smaller Allan variance. Thus, what has been shown for a limited set of simulated signals in Fig. 3 is now backed up with analytical calculations. It contradicts the very physical sense of the Allan variance because with R>1 no averaging and no oversampling means smaller noise, which just is meaningless. In order to figure out whether it is possible to use the modified Allan deviation plot in some another way we can use numerical experiments.

Approbation results

In the case of a magnetometer, a reasonable approach to judging of the accuracy of *X*-, *Y*- and *Z*-measurements is calculation of the magnetic flux density length:

$$B = \sqrt{B_x^2 + B_y^2 + B_z^2} \tag{7}$$

where B_x, B_y and B_z are projections of the magnetic flux density vector onto the axes *X*, *Y* and *Z* correspondingly of some Cartesian coordinate system assigned with the magnetometer. Since *X*-, *Y*- and *Z*-readings are produced in the magnetometer frame, not the earth frame, it is not possible to consider them independently without additional tools of measuring the magnetometer attitude.

In the absence of interferences, the vector length is expected to be invariant when a magnetometer is rotated in a homogeneous magnetic field. We make use of this invariance in order to evaluate the applicability of the Allan deviation plot for the same sensor with different sample frequencies applied. For handling noise, one should first decide on the trade-off between the allowable noise level and a time period one can afford to spend on obtaining extra sensor readings in order to average them and eliminate noise. Once the choice has been done, i.e. a point (*X*,*Y*) is selected on the Allan deviation plot, its *Y* value has the following meaning. Approximately 68 % of measurements will have the deviation from the true noise level, numerically equal to *Y*. Thus, one subtracts *Y* from the measurements and partly eliminates noise in this way. If one rotates a magnetometer, takes triples of values $\left(B_x,\ B_y,\ B_z\right)$, eliminates noise, then calculates the magnetic flux density by (7), stores the calculated values of the magnetic flux density in a vector and computes the variance of the vector items, the variance tends to zero if noise has been properly eliminated. Hence, (7) can be used as a measure of how well noise has been filtered out.

In order to verify whether the Allan deviation plot computed for signals sampled with a frequency exceeding the data refreshment rate is applicable we take the following steps.

1. The first step is to build the Allan deviation plots for a triaxial magnetometer. Keeping the magnetometer steady, we acquire three signals for *X*, *Y*, and *Z* axes correspondingly, with some initial sample frequency f_1. The time series should be long enough for the Allan deviation plot to be meaningful (no less than 16000 points as has been said previously).

2. With a fixed step Δm for the averaging factor *m*, starting from $m = 1$, we pick *Y*-values of the corresponding *X*-coordinates values $\tau = m\tau_0$ on the Allan deviation plot for axis *X*. Besides, we consider the maximum *Y*-value of the Allan deviation plot, whether it corresponds to either of $(1+ k\Delta m)\ \tau_0$ or not (*k* is 1,

2, 3,….). All the selected *Y*-values should be stored in the *i*-th row of matrix AllanDeviationsX, where number *i* is the number of an experiment (the number of the sample frequency being considered). The same procedure should be done with two other plots, for axes *Y* and *Z* (two additional matrices are needed, AllanDeviationsY and AllanDeviationsZ).

3. Steps 1 and 2 are repeated for sample frequencies $f_2 = 2f_1$, $f_3 = 3f_1$, etc. *Y*-values for the newly obtained Allan deviation plots should be stored in the corresponding rows of matrices AllanDeviationsX, AllanDeviationsY and AllanDeviationsZ. The Allan deviation characterizes a sensor, not data it produces. Therefore, the matrices can be used later for eliminating noise from new signals.

4. We lift the magnetometer from a nearly flat surface, which can be a table for example, rotate it in air in the plane more or less parallel to the surface and put it down on the surface again without much translational movement. In this way two out of the three projections of the magnetic flux density vector will be affected much more than the third one. However, firstly, changes to two projections are sufficient, and secondly, we do not need additional mechanisms to fix the device – laid on the table, it will be kept steady for a while.

5. Keeping the device static in an arbitrary (and unknown) position, we obtain measurements for all the three axes and store them into three vectors, Bx, By and Bz, correspondingly. The process should be repeated for a number of arbitrary rotations.

6. The first item of matrix AllanDeviationsX (AllanDeviationsX[0][0]) should be subtracted from vector Bx. Similarly, the first items of two other matrices are subtracted from two other vectors correspondingly. Then we calculate the magnetic flux density using (7) for each triple Bx[j], By[j] and Bz[j] and the variance of the obtained lengths of the magnetic flux density vector. The variance is stored in a new matrix, Results.

7. The previous step is repeated for each item in matrices AllanDeviationsX, AllanDeviationsY and AllanDeviationsZ. The variance should be stored in Results, in the cell with the same indices.

8. We compare the items of matrix Results. Items close to zero indicate proper noise elimination.

By taking the described steps we detected that the Allan deviation plots are completely inapplicable if signals are sampled more frequently than new data can be available (no matter which *Y*-values are considered, minimum or maximum). The fact that for $m = 1$ one observes smaller *Y*-values in these plots than in the Allan deviation plot for data polled as frequently as new data appear turned out to be misleading. The least variance in the magnetic flux density was observed when a signal contained no systematic repetitions.

The described verification approach is not flawless. However, we can ignore its imperfections (assumptions about a homogeneous magnetic field and the absence of distortions) since they influence all measurement results in the same way.

Conclusions

The work considers a problem of interpreting Allan deviation plots dissimilar to the classic textbook case. Being unable to interpret the Allan deviation plot properly often means a failure to cope with noise terms inevitably present in inertial sensors. The problem arose from practice, when improperly sampled signals from a magnetometer caused the Allan deviation plot far dissimilar to the classical case and thus difficult to decipher.

We have found relationships between the sample frequency of signals and the shape of the Allan deviation plot for these signals, using both simulation techniques and analytical considerations. Using the fact that the magnetic flux density vector length remains invariant when rotating a magnetometer in a uniform field, we have experimentally found that Allan deviation plots for signals sampled with a frequency exceeding the data refreshment rate are not applicable at all. One needs to remove repeated measurement points before applying the Allan variance method.

References

1. Lee, J., Mellifont, R., & Burkett, B. (2010). The use of a single inertial sensor to identify stride, step, and stance durations of running gait. *Journal of Science and Medicine in Sport, 13(2)*, 270–273. doi:10.1016/j.jsams.2009.01.005

2. Tang, Z., Sekine, M., Tamura, T., Tanaka, N., Yoshida, M., & Chen, W. (2015). Measurement and Estimation of 3D Orientation using Magnetic and Inertial Sensors. *Advanced Biomedical Engineering, 4*, 135–143. doi:10.14326/abe.4.135

3. Mao, A., Ma, X., He, Y., & Luo, J. (2017). Highly Portable, Sensor-Based System for Human Fall Monitoring. *Sensors, 17(9)*, 2096. doi:10.3390/s17092096

4. Marina, H., Pereda, F., Giron-Sierra, J., & Espinosa, F. (2012). UAV attitude estimation using unscented Kalman filter and TRIAD. *IEEE Transactions on Industrial Electronics, 59(11)*, 4465–4474. Retrieved from https://arxiv.org/pdf/1609.07436

5. Faragher, R. (2012). Understanding the Basis of the Kalman Filter Via a Simple and Intuitive Derivation. *IEEE Signal Processing Magazine, 29(5)*, 128–132. doi: 10.1109/MSP.2012.2203621.

6. Mahony, R., Hamel, T., & Pimlin, J.-M. (2008). Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control, 53(5)*, 1203–1218.

7. El-Sheimy, N., Hou, H., & Niu, X. (2008). Analysis and Modeling of Inertial Sensors Using Allan Variance. *IEEE Transactions on Instrumentation and Measurement, 57(1)*, 140–149. doi: 10.1109/TIM.2007.9086354.

8. Vukmirica, V., Trajkovski, I., & Asanović, N. (2010). Two Methods for the Determination of Inertial Sensor Parameters. *Scientific Technical Review, 60(3–4)*, 27–33.

9. U.S. Army Research, Development and Engineering Center. (2015). Allan variance calculation for nonuniformly spaced input data (Publication No. ARWSE-TR-14011). Retrieved from https://apps.dtic.mil/dtic/tr/fulltext/u2/a616850.pdf

10. Friederichs, T. (2019). Analysis of geodetic time series using Allan variances. Retrieved from https://elib.uni-stuttgart.de/bitstream/11682/3866/1/Friederichs.pdf

11. University of Cambridge. Computer Laboratory. (2007). An introduction to inertial navigation (Publication No. UCAM-CL-TR-696). Retrieved from https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-696.pdf

12. Riley, W. J. (2008). *Handbook of frequency stability analysis*. Washington: U. S. Government Printing Office.

13. Barrett J. M. (2014). *Analyzing and modeling low-cost MEMS IMUs for use in an inertial navigation system* (Master's thesis). Retrieved from https://web.wpi.edu/Pubs/ETD/Available/etd-043014-163543/unrestricted/jbarrettMSThesis.pdf