

Algorithms for frequency jump detection

W J Riley

Hamilton Technical Services, 650 Distant Island Drive, Beaufort, SC 29907, USA

E-mail: bill@wriley.com

Received 19 June 2008, in final form 11 September 2008

Published 5 December 2008

Online at stacks.iop.org/Met/45/S154

Abstract

The detection of jumps in a frequency record is a challenging problem by either visual or mathematical means. The former takes considerable experience and judgment, and is therefore quite subjective, but has the advantage of providing insight into device behaviour. The latter is more impartial and consistent, and can be automated. In combination, mathematical jump detection can be applied for automatic clock testing and monitoring. If a jump is detected, the record can then be inspected visually before deciding on the action required. This paper describes frequency jump detection algorithms that are included in a program for frequency stability analysis.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

There is a large body of technical literature on the subject of detecting discontinuities in time series (often referred to as regime shift or change-point analysis) [1–3]. Many such algorithms involve comparing the mean values of two adjacent windows that are moved through the data, while others use cumulative sum charts and other similar techniques. Some of these methods require a data model whose parameters must be determined, while other methods are non-parametric. They can variously detect outliers, transients, level shifts and changes in variance for data having different noise characteristics and trends [4]. The emphasis here is on detecting jumps in reasonably white frequency residuals after outliers and deterministic drift have been removed from the data. Robust and effective means exist for outlier removal based on the median absolute deviation (MAD), several models exist for fitting and removing frequency drift and the dynamic Allan deviation (DAVAR) can be used to examine clock data for stability changes [5]. This paper describes several simple and intuitive techniques for detecting, analysing and displaying frequency jumps.

2. Frequency jump detection algorithms

This paper describes two methods for frequency jump detection, using either block averaging (BLKAVG) or sequential averaging (SEQAVG) of the frequency data. It also

covers the classic cumulative sum (CUMSUM) method for analysing data for changes in mean [2].

The BLKAVG algorithm is very intuitive. It simply compares the average values within two non-overlapping moving analysis windows, declaring a jump if their difference exceeds a certain limit that remains the same throughout the data set. This detection method has been used by the author for many years to automatically screen clock data for frequency jumps.

The SEQAVG algorithm is based on the Rodionov Sequential t -Test Analysis of Regime Shifts (STARS) algorithm [6]. It does not use fixed analysis windows, but rather scans the data sequentially. If a suspect point is found that exceeds the jump threshold, the next block of data is examined to confirm or reject the jump. Instead of using a jump threshold based on the Student's t -test [7], the SEQAVG method uses a fractional frequency jump limit.

The CUMSUM method is well explained by Taylor in [2]. It is effective for locating and quantifying single jumps in otherwise well-behaved frequency data. Interestingly, it is identical to the corresponding normalized (mean-removed) phase data, where a sharp slope change denotes a frequency jump.

2.1. Block average algorithm

The Block Average (BLKAVG) frequency jump detection algorithm compares the difference between the average values within a pair of adjacent windows of adjustable length against

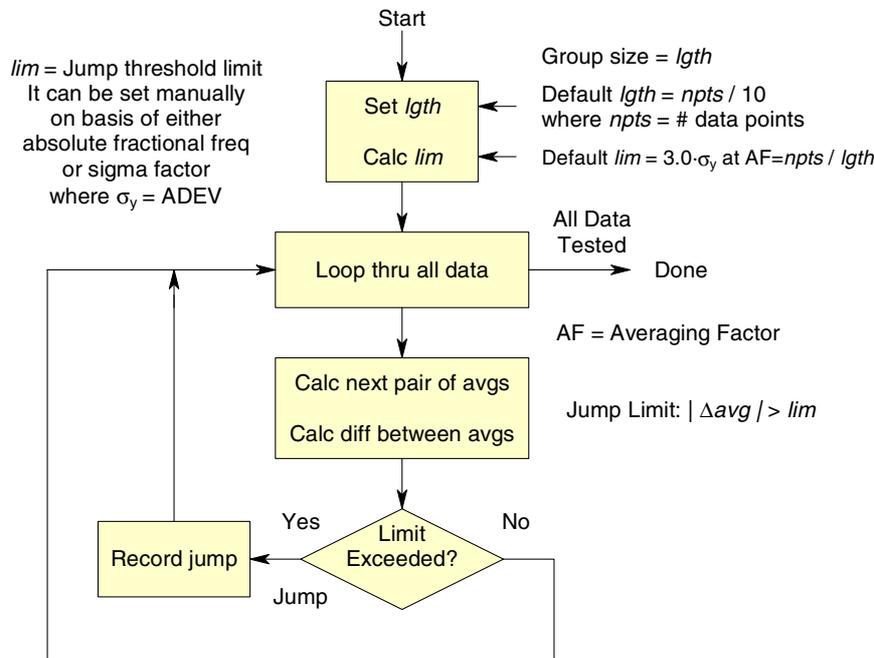


Figure 1. Block average jump detection flowchart.

a selectable jump threshold. That threshold can be either an absolute fractional frequency limit or a sigma factor based on the Allan deviation at an averaging factor equal to the window length over the entire data set. The start of the data windows can be offset up to the amount that the data that exceeds the total span of the data windows. The BLKAVG jump detection algorithm is shown in the flowchart of figure 1.

The visual presentation of a BLKAVG frequency jump analysis can be enhanced by plotting the average value of each analysis window along with the underlying frequency data.

2.2. Sequential average algorithm

The Sequential Average (SEQAVG) method is similar, but, instead of dividing the frequency record into fixed blocks, it searches the data sequentially. If a potential jump is detected, the following points are examined in an averaging window to accept or reject the jump. The SEQAVG plot shows the actual frequency averages in each jump regime, and is able to locate the jump more closely, but does not show trends in the data. The SEQAVG jump detection algorithm is shown in the flowchart of figure 2.

The visual presentation of a SEQAVG frequency jump analysis can also be enhanced by plotting the average value of each jump regime along with the underlying frequency data.

A comparison between the BLKAVG and SEQAVG methods is shown in table 1.

2.3. Cumulative sum jump detection

The Cumulative Sum (CUSUM) is a classic change-point analysis technique that uses the cumulative sum of the differences between the current value and the overall

average [2]. Starting with $S_0 = 0$,

$$S_i = S_{i-1} + (y_i - \bar{y}) \quad (1)$$

where S_i is the cumulative sum and \bar{y} is the average of the data.

The CUSUM slope indicates the value of the data with respect to the overall average. A flat cumulative sum indicates that the data is near the average value, a straight cumulative sum indicates a period of constant data values, and, most importantly, a sudden change in the CUSUM slope indicates a jump in the data. The CUSUM plot for a data set having a single jump will have a V or inverted V shape. The value of the jump can be determined by the sum of the two slopes of the CUSUM curve. For N data points having a maximum or minimum cumulative sum value M at point P , the frequency jump is equal to

$$\frac{M}{P-1} + \frac{M}{N-P}. \quad (2)$$

2.4. STARS algorithm

Recent progress in detecting 'regime shifts' in ecological records has resulted in a jump detection algorithm called the Sequential t -Test Analysis of Regime Shifts (STARS), which could be useful for analysing clock frequency records [6]. The STARS algorithm uses the Student's t -test [7] to compare the mean values of two analysis windows that are moved sequentially through the data record, searching for differences that are associated with a significant jump. The problem with the STARS method is that it is too sensitive. Frequency stability mainly involves the analysis of noise. The t -test indicates whether, at a certain confidence level, there has been a shift in the mean, generally for small sample sizes. For the analysis of noise, this detects too many changes. The better criterion for noise is either an absolute frequency deviation, or a jump that exceeds the Allan deviation by a certain factor.

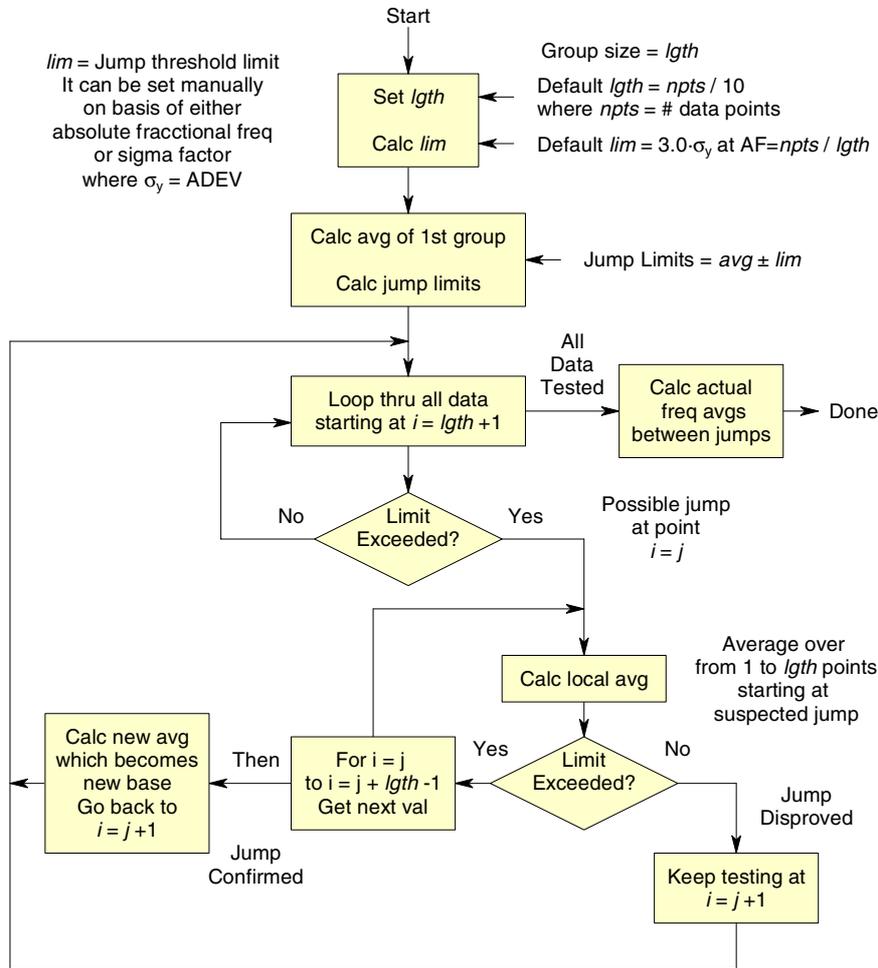


Figure 2. Sequential average (SEQAVG) jump detection algorithm flowchart.

Table 1. Comparison between BLKAVG and SEQAVG frequency jump detection methods.

Method	Advantages	Disadvantages
Block Average (BLKAVG)	<ol style="list-style-type: none"> 1. Simple, intuitive, easy to understand, obvious from plot. 2. Shows non-jump changes. 3. Handles drift better, shows it as block averages, not jumps. 	<ol style="list-style-type: none"> 1. Does not get exact jump location. 2. Does not get actual frequency averages. 3. May not get accurate jump size.
Sequential Average (SEQAVG)	<ol style="list-style-type: none"> 1. Get right jump location. 2. Get actual frequency averages within jump regimes. 3. Cleaner plot. 4. Better jump size estimates. 5. No offset parameter needed. 6. Window length less critical (if not too long). 	<ol style="list-style-type: none"> 1. More complex algorithm. 2. Does not show non-jump changes. 3. Depends on average at beginning of each regime. 4. Shows drift as staircase of jumps. 5. Slower (but not significantly).

3. Frequency jump detection examples

Examples of the BLKAVG, SEQAVG and CUSUM frequency jump detection methods are shown below for a set of 1024 points of simulated frequency data having white FM noise at a level of 1×10^{-12} at a 1 s sampling time plus a step of -1×10^{-12} at the centre of the record. The objective is to correctly detect this frequency jump in the presence of an equal amount of noise, and quantify both the location of the jump

and its value, preferably using default detection parameters (an analysis window size based on the record length and a frequency jump threshold equal to a multiple of the Allan deviation).

3.1. BLKAVG jump detection

This example of the BLKAVG frequency jump detection algorithm is shown in figure 3. The light horizontal lines on

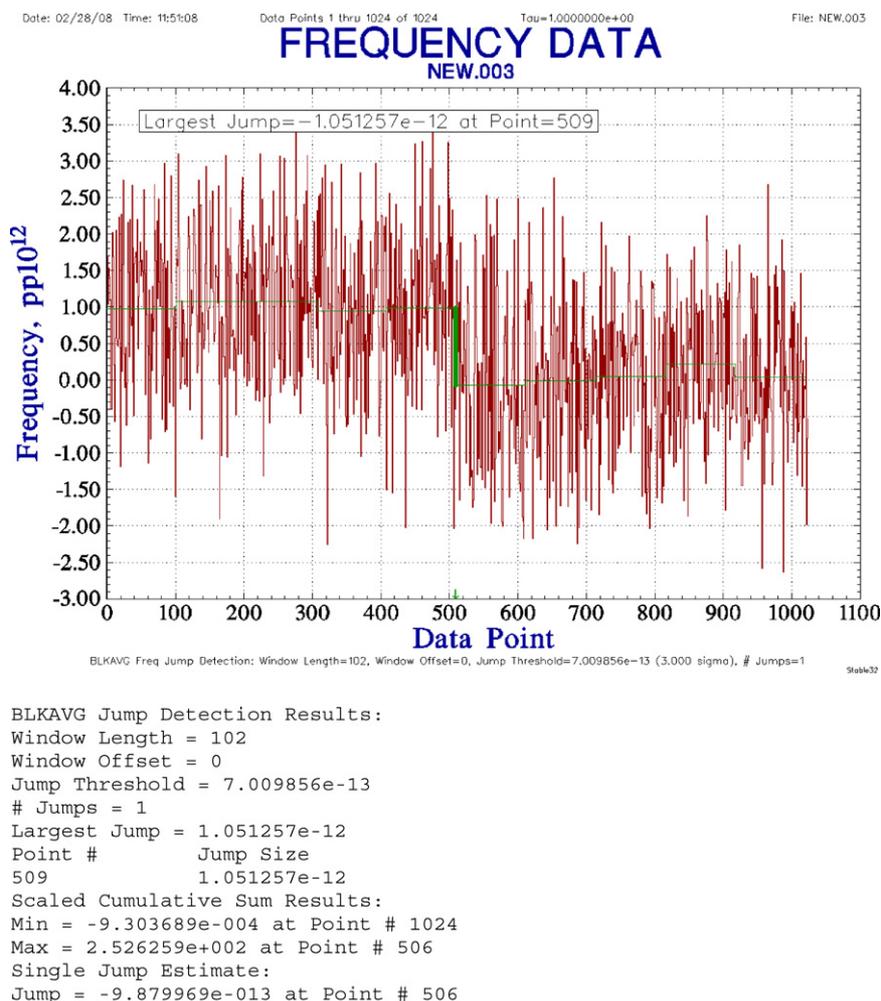


Figure 3. BLKAVG plot and results.

the plot are the average values of the frequency data over the analysis windows, and a detected frequency jump is shown by a heavy vertical line.

The default window length is one-tenth of the record length or five, whichever is larger. The default window offset is zero, and it can have a maximum value equal to the number of frequency data points modulo the window length. The default jump threshold is three times the value of the Allan deviation of the entire data set at an averaging factor equal to the window length. The jump detection behaviour can be optimized for the particular application by adjusting the detection parameters.

3.2. SEQAVG jump detection

An example of the SEQAVG frequency jump detection algorithm for the same data set is shown in figure 4. Here the light horizontal lines on the plot are the average values of the frequency data over the entire jump regimes, and the detected frequency jump is again shown by the heavy vertical line near the centre of the plot.

The default window length and jump threshold are the same as for the BLKAVG method, and can also be tuned for the particular application. The window offset parameter does not

apply to the SEQAVG method. One jump is correctly detected for window lengths between 28 and 220 for these data.

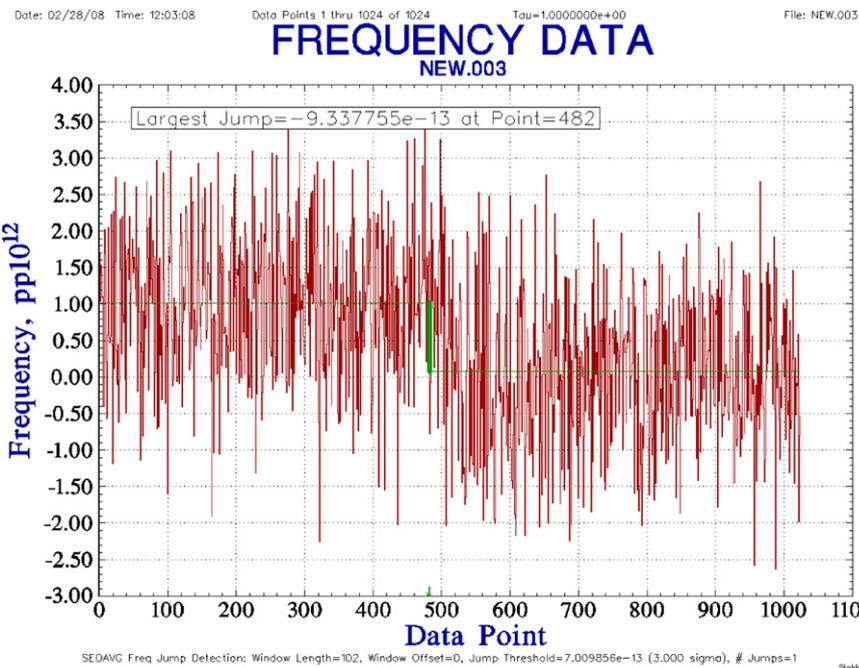
3.3. SEQAVG jump location

The SEQAVG frequency jump detection algorithm tends to estimate the jump location earlier than it actually is because it reports it at the beginning of the averaging window in which the jump is confirmed. That bias can be reduced significantly by averaging the jump location estimates for the normal forward and reversed data sets. That improved jump location estimate is given by

$$J = \frac{F + (N - R)}{2} \quad (3)$$

where J is the estimated jump location, F is the forward jump location, R is the reverse jump location and N is the number of data points.

For example, the forward and reverse default SEQAVG jump detection results for the same 1024-point set of simulated frequency data are 482 and 478, respectively. The combined jump location estimate $J = (482 + 1024 - 478)/2 = 514$ is very close to that expected, 512, much closer than either the forward or reverse estimates alone.



```

SEQAVG Jump Detection Results:
Window Length = 102
Jump Threshold = 7.009856e-13
# Jumps = 1
Largest Jump = -9.337755e-13
Point #      Jump Size
482          9.337755e-13
Scaled Cumulative Sum Results:
Min = -9.303689e-004 at Point # 1024
Max = 2.526259e+002 at Point # 506
Single Jump Estimate:
Jump = -9.879969e-013 at Point # 506

```

Figure 4. SEQAVG plot and results.

3.4. SEQAVG jump magnitude

The jump magnitudes are similar in each direction (differing only by about 3.5%), and have an average value of 9.18×10^{-13} , reasonably close to the nominal value of 1×10^{-12} and the actual difference between the two halves of the data, 9.77×10^{-13} . Reducing the window size to 50 yields an average jump magnitude of 9.54×10^{-13} , 2.4% below the actual.

3.5. Summary of BLKAVG and SEQAVG results

The previous examples of frequency jump detection methods showed the analysis an abrupt frequency change of 1×10^{-12} at the centre of a 1024-point frequency record of simulated white FM noise having an Allan deviation of 1×10^{-12} at the 1 s sampling interval. The jump was correctly identified using the default BLKAVG detection parameters as having a magnitude of 1.05×10^{-12} at point 509. The calculated 1 s overlapping Allan deviation of the entire frequency record, 0.992×10^{-12} , is not significantly affected by the jump, which causes the stability plot to flatten and then go upward at longer averaging times. Varying the window offset over its allowable range of 0 to 24 had little effect on the magnitude of the detected frequency jump, which varied from 0.993 to 1.116×10^{-12} . The frequency jump is visually apparent in the

data plot, especially if they are averaged by a factor of about ten, and the detection algorithms are able to quantify the jump automatically at least as well as can be done manually.

3.6. CUMSUM plot

The BLKAVG and SEQAVG frequency jump detection methods can also make use of cumulative sum (CUSUM) calculations to help identify and estimate the location and value of frequency jumps, as shown in figure 5.

Inspection of the CUSUM plot provides a good indication of the confidence in the jump determination and can show the presence of other smaller, below-threshold jumps. For a single dominant jump, it does an excellent job of locating and quantifying the jump. For frequency data, the CUSUM is the same as its corresponding phase data after frequency normalization.

4. Confidence factor

A nice aspect of the CUSUM method is its ability to provide an estimate of the confidence factor for a detected jump. This can be accomplished by randomly reordering the frequency data and calculating the CUSUM for many such cases,

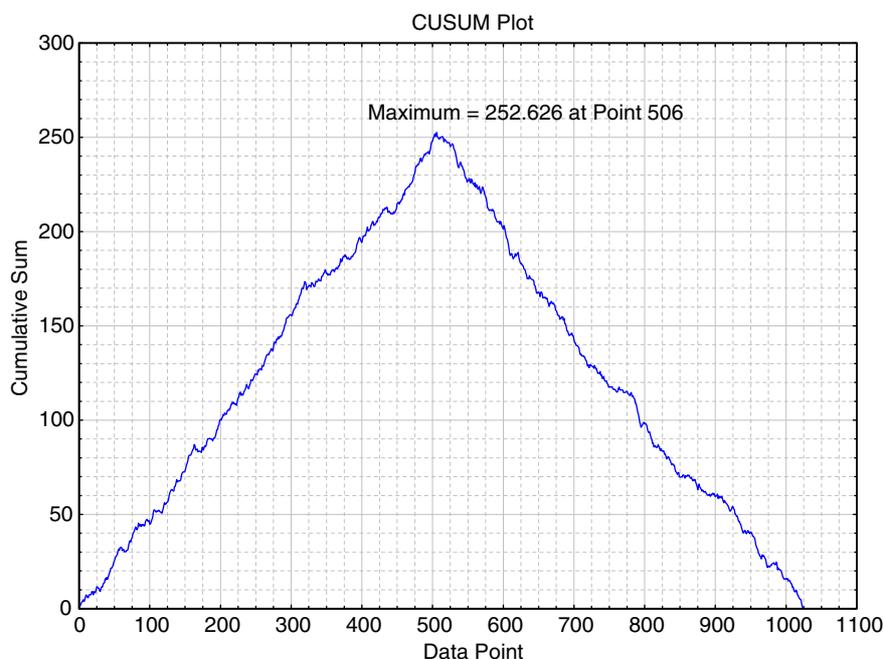


Figure 5. CUSUM plot.

thereby establishing a basis of comparison for the actual jump estimation. For example, if in 995 out of 1000 such cases, the CUSUM range is less than that for the actual data, the confidence in the jump is 99.5%. In most cases for data set sizes and thresholds appropriate for clock frequency data, the confidence level is above 99.9% when a jump is detected.

5. Jump detection criteria

Important considerations are the ability to reliably detect jumps of a specified magnitude with minimal false alarms in the presence of common clock noises and slow frequency trends. Frequency jumps can be the most significant limitation to clock performance in many applications. While there is no substitute for human judgment in assessing clock behaviour, automated jump detection can be a valuable supplement to that, especially in a large-scale production situation (e.g. commercial telecom clocks), or for critical applications (e.g. GPS or Galileo satellite clocks).

6. Jump display and reporting

Frequency jump detection is naturally associated with a frequency data plot as shown in the preceding examples. The frequency averages are shown as green horizontal lines extending over the analysis data windows. Those average values are connected by vertical steps that, in the case of a detected frequency jump, are shown as heavy green lines. The jump detection parameters are shown as an annotation at the bottom of the plot, the largest jump magnitude and location (if any) is shown as an optional message, and a CUSUM plot inset may also be displayed. Complete information about the jump detection parameters and results is written to the Windows clipboard, and can be pasted into a plot note or

another application. The frequency averages and cumulative sums corresponding to each data point are automatically written to data files. The colour of the frequency data in the BLKAVG and SEQAVG plots can be changed to a faint yellow so that the jump analysis can be better seen.

7. Jump detection limitations

The most significant limitation of the BLKAVG frequency jump detection algorithm is that it uses analysis windows of a fixed (although selectable) size. That means that jumps can be detected only at certain places in the frequency record, and, accordingly, the magnitude of the detected jumps are not necessarily those of the underlying data but simply the change in the average value at the boundary of the two analysis windows. This limitation can be reduced by changing the size of the analysis window and/or their offset. The window size is a tradeoff between better time resolution (short) and better noise averaging (long). In contrast, the SEQAVG and CUSUM methods can detect a jump at any point in the data record. The SEQAVG jump location can be improved by analysing the data in both directions. All of these jump detection methods are sensitive to drift and divergent noise.

8. Prewhitening

It can be hard to tell the difference between divergent noise and an actual jump. For example, random walk FM noise with no jump resembles white FM noise having an actual jump in mean value. Jump detection can sometimes be improved by removing not only outliers and deterministic drift, but also by removing flicker ('pink') and random walk ('red') noise from the data [8]. That prewhitening process can help to distinguish between true jumps and the lurching behaviour caused by

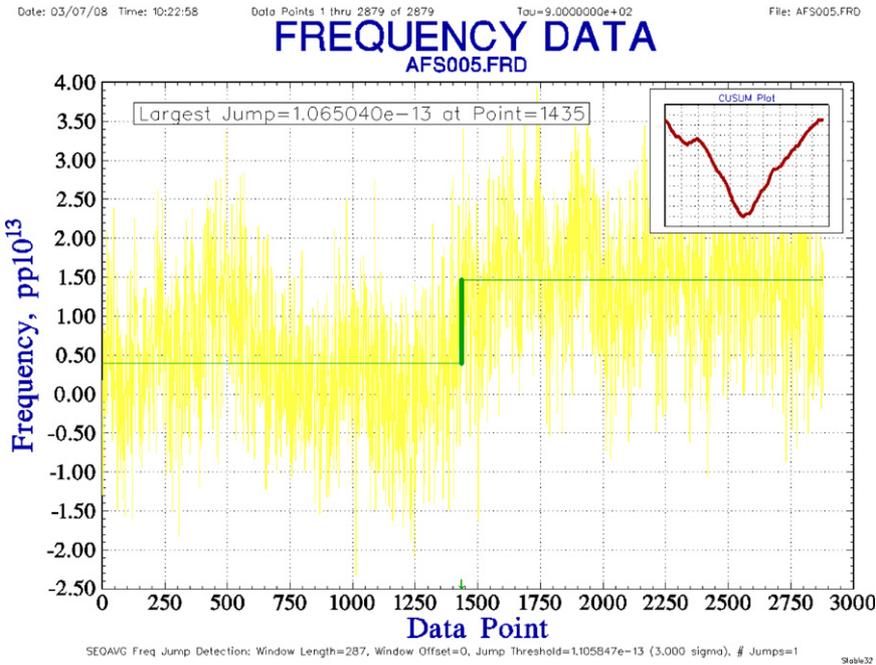


Figure 6. SEQAVG frequency jump detection with simulated white and flicker FM noise.

divergent noise. Such noise removal, although inexact, can be accomplished by determining the lag 1 autocorrelation coefficient, ρ_1 , of the data and using that to remove an AR(1) autoregressive noise component [9, 10]:

$$y(t) = y(t+1) - \rho_1 \cdot y(t). \quad (4)$$

A ρ_1 value of 0 corresponds to white noise, while increasingly positive ρ_1 values result from pink and red noise spectra, with a value of 1 associated with random walk noise. Examples of AR(1) noise processes and their autocorrelations and spectra are shown in [10]. The lurches of the more divergent noise processes can often resemble jumps.

Consider, for example, a sample of random walk noise having a $\rho_1 = 0.989$ and a power law exponent $\alpha = -1.90$ at an averaging factor of 1. If that noise data is prewhitened as described above, the resulting residuals are white ($\rho_1 = -0.041$, $\alpha = +0.09$) and are without any apparent jump.

A determination of ρ_1 can itself be affected by a jump, and techniques have been described for sectioning the data and correcting for bias [8]. The prewhitening operation will reduce the size of the jump, but the process is nevertheless reasonably effective in distinguishing between divergent noise and an actual jump, whose location is preserved.

9. Additional examples

Another example of SEQAVG frequency jump detection is shown in figure 6, again using the default detection parameters plus a CUSUM plot inset. This is a 30-day set of 2880 points of $\tau = 900$ s (15 min) simulated clock data having $2 \times 10^{-12} \tau^{-1/2}$ white FM noise and flicker FM noise at a level of 2×10^{-14} plus a jump of $+1 \times 10^{-13}$ at the centre (point 1440).

The jump is detected, located and quantified with reasonable accuracy. The CUSUM analysis characterizes the jump as having a value of $+1.06 \times 10^{-13}$ at point 1435. The positive frequency ‘lurch’ near point 500 is not a jump but rather a result of the divergent flicker FM noise. It is visible in both the frequency and CUSUM plots, and its average frequency change is about half the threshold level. The same simulated noise without the frequency jump has a CUSUM range of 232, about one-third the jump size and no jumps are found with the same detection parameters. A typical set of 1000 random CUSUMs has a range of 118, a mean of 82 and a standard deviation of 19.3, so the actual CUSUM range of 767 is many standard deviations larger, implying a high confidence in the jump detection.

Other analysis methods such as Allan deviation and autocorrelation function plots can also provide insight into the presence and effect of a frequency jump. An upward slope of a drift-removed Allan deviation plot at long τ indicates a significant frequency disturbance of some sort since flicker FM noise would have a flat characteristic. Although usually associated with random walk FM noise, a $\tau^{1/2}$ slope can also be caused by a step in average frequency. Similarly, a long positive autocorrelation tail at large lags is a sign of either divergent noise or some sort of long-term frequency change. However, neither of these plots can be uniquely or simply related to a frequency jump.

Prewhitening by removing an AR(1) autoregression process from a mixture of flicker and random walk FM noise results in much whiter frequency data. Even though any frequency jump in the data would be reduced in magnitude by the prewhitening, simulations have shown that a jump can still be detected at the right location with a BLKAVG or SEQAVG analysis using a 3-sigma threshold. This indicates that it is actually a jump and not simply the result of the underlying divergent noise.

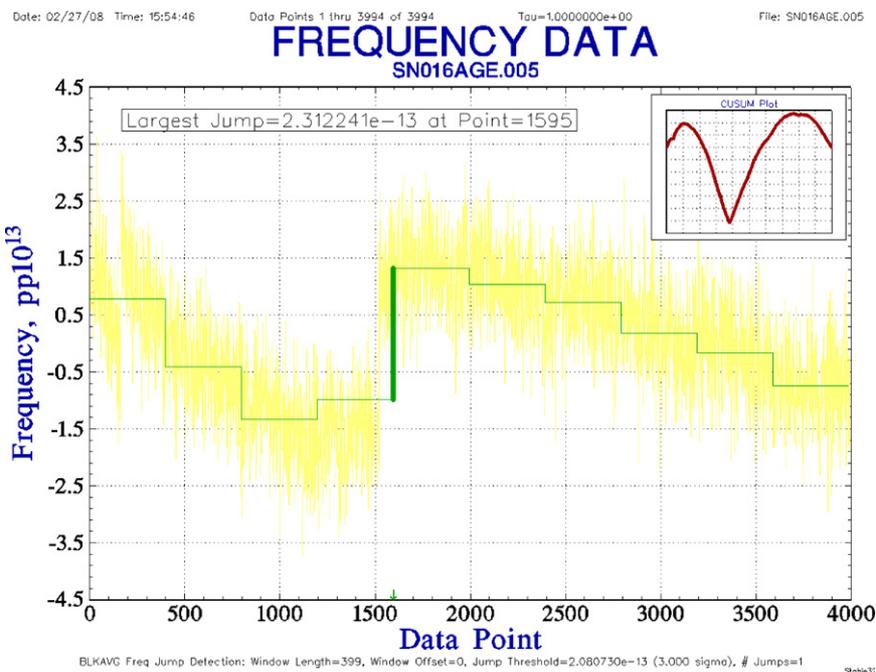


Figure 7. BLKAVG frequency jump detection for RFS stabilization record.

These frequency jump detection and analysis methods have also been applied to actual clock data, including the 42-day, $\tau = 900$ s early stabilization frequency residual record of a rubidium frequency standard after removal of a log fit shown in figure 7.

This record shows two frequency jumps, the first during the early rapid stabilization and a second more prominent jump at about point 1600. The relatively long analysis window size (399) delays detection of the large jump and the default detection threshold ignores the first smaller jump. The CUSUM inset shows only one sharp change because the first jump is actually more like a transient, and the multiple peaks distort its jump magnitude determination. If the window size is changed to 100, the analysis still detects a single jump, whose location is then determined more exactly while still providing adequate noise averaging. Reducing the jump threshold to 5×10^{-14} produces three detected jumps (including the first one) but splits the second jump into two parts. Clearly, user judgment is needed when performing a jump analysis on a complex record.

In another similar example, the frequency residuals show a single large lurch and several smaller jumps. The default 3-sigma BLKAVG jump analysis works well for this slower frequency discontinuity. The CUSUM plot has its sharpest and largest peak at the large lurch, but is too complex to provide a meaningful jump estimate. Multiple smaller jumps can be detected by reducing the window size and threshold.

10. Conclusions

This paper has described several methods for the detection and analysis of frequency jumps in clock data. The BLKAVG and SEQAVG algorithms, along with a CUSUM plot, are reasonably consistent and effective ways to examine such

jumps. Those methods have been successfully implemented into a program for frequency stability analysis [11].

References

- [1] Rodionov S N 2005 A brief overview of the regime shift detection methods *Large-Scale Disturbances (Regime Shifts) and Recovery in Aquatic Ecosystems: Challenges for Management Toward Sustainability* ed V Velikova and N Chipev *UNESCO-ROSTE/BAS Workshop on Regime Shifts (Varna, Bulgaria, 14–16 June 2005)* pp 17–24 http://www.beringclimate.noaa.gov/regimes/rodionov_overview.pdf
- [2] Taylor W A 2000 Change-point analysis: a powerful new tool for detecting changes <http://www.variation.com/cpa/tech/changepoint.html>
- [3] Basseville M and Nikiforov I V 1993 *Detection of Abrupt Changes—Theory and Application* (Englewood Cliffs, NJ: Prentice-Hall) ISBN 0-13-126780-9 <http://www.irisa.fr/sisthem/kniga/kniga.pdf>
- [4] Tsay R S 1988 Outliers, level shifts and variance changes in time series *J. Forecast.* **7** 1–20
- [5] Riley W J 2007 *Handbook of Frequency Stability Analysis* (Beaufort, SC: Hamilton Technical Services)
- [6] Rodionov S N 2004 A sequential algorithm for testing climate regime shifts *Geophys. Res. Lett.* **31** L09204
- [7] Student's *t*-test, *Wikipedia* http://en.wikipedia.org/wiki/Student%27s_t-test
- [8] Rodionov S N 2006 Use of prewhitening in climate regime shift detection *Geophys. Res. Lett.* **33** L12707
- [9] Autoregressive Moving Average Model, *Wikipedia* http://en.wikipedia.org/wiki/Autoregressive_moving_average_model
- [10] Wagenmakers E J, Farrell S and Ratcliff R 2004 Estimation and interpretation of $1/f^\alpha$ noise in human cognition *Psychonomic Bull. Rev.* **11** 579–615 <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1479451>
- [11] *Stable32*, Program for Frequency Stability Analysis (Beaufort, SC: Hamilton Technical Services) <http://www.wriley.com/Brochure.pdf>