PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Delayed comparison error minimization for frequency domain state estimation in structures subjected to high-rate boundary change

Scheppegrell, James, Downey, Austin, Moura, Adriane, Dodson, Jacob

James Scheppegrell, Austin Downey, Adriane G. Moura, Jacob Dodson, "Delayed comparison error minimization for frequency domain state estimation in structures subjected to high-rate boundary change," Proc. SPIE 11593, Health Monitoring of Structural and Biological Systems XV, 115932Q (22 March 2021); doi: 10.1117/12.2583331

SPIE.

Event: SPIE Smart Structures + Nondestructive Evaluation, 2021, Online Only

Delayed comparison error minimization for frequency domain state estimation in structures subjected to high-rate boundary change

James Scheppegrell[a,b], Austin Downey[c,d], Adriane G. Moura[a,b], and Jacob Dodson[a]

[a] Air Force Research Laboratory , Eglin Air Force Base , 203 Eglin Blvd, Eglin AFB, FL 32542,
[b] Applied Research Associates, Emerald Coast Division, 956 John Sims Pkwy W., Niceville, FL 32578,
[c]Department of Mechanical Engineering, University of South Carolina, Columbia, SC, USA
[d]Department of Civil and Environmental Engineering, University of South Carolina, Columbia SC, USA

ABSTRACT

Many structural systems, such as aircraft, orbital infrastructure, and energy harvesting devices, experience dynamic forces along with changing structural boundary conditions. Collecting and analyzing data on these systems provides useful insight that aids design, evaluation, and function. For real-time decision-making on systems experiencing high-rate changes, completing assessments quickly enough to be relevant poses a unique set of challenges. In systems sufficiently understood and well defined, determining a system's state that experiences high-rate structural boundary condition changes can be accomplished by monitoring its frequency response. In this work, methods of frequency detection applicable to real-time state estimation of structures experiencing high-rate boundary changes were investigated; progress and findings in extracting the frequency response of a structure in real-time are presented here. A novel Delayed Comparison Error Minimization technique is presented and experimentally validated using the DROPBEAR experimental testbed at the Air Force Research Laboratory. This testbench consists of an oscillating beam with one end fixed and roller support that can move along the beam's length. Real-time estimation of pin location through the measurement of beam motion was performed using the novel Delayed Comparison Error Minimization technique. Results are compared against an FFT-based method with a variety of window lengths. The latency and precision of this method are analyzed, and the results from each method are compared, with a discussion on the applicability of each method.

Keywords: Pitch Detection, Structural Health Monitoring, Structural Dynamics, State Estimation

1. INTRODUCTION

There are many systems and structures in use and under development that experience high-rate dynamic events, defined as changes occurring on a time scale of under 100 ms.[1] Many of the parameters being monitored in these systems are parts of processes being managed continuously by control loops, and therefore part of the normal operation of the system. For some systems, there exists a possibility that they will be subjected to conditions or states which are undesirable or can result in damage. Anticipating such a condition by detecting and reacting to an intermediate state before it progresses further can result in an improved outcome for the system. These applications have helped create a desire for an observer which can assess the state of such structures and systems, determine the conditions, and make decisions based on the state quickly enough that the assessment has not become irrelevant. Such an observer would find potential applications across disparate fields, including machinery and automation, blast mitigation, and hypersonic aircraft.[2–5] While the applications may not appear to have much in common, the commonalities in their needs have created motivation to develop general-purpose tools that can be used in creating the solutions needed for each case.

The goal of this paper is to design, implement, and demonstrate a frequency-based observer capable of estimating the state of a complex system with the necessary precision and speed. To achieve this goal, the design of a useful observer must take into consideration the structure's properties, collect the most useful information

Health Monitoring of Structural and Biological Systems XV, edited by Paul Fromme, Zhongqing Su, Proc. of SPIE Vol. 11593, 115932Q · © 2021 SPIE
CCC code: 0277-786X/21/$21 · doi: 10.1117/12.2583331

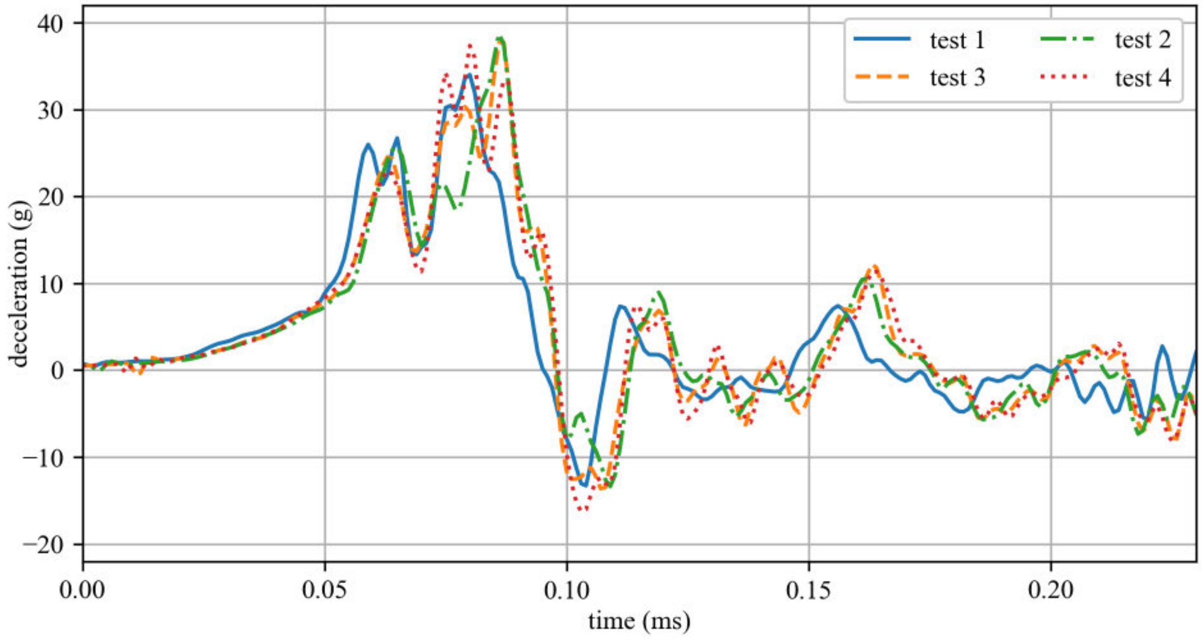

Figure 1. Data from a sub-second system showing response from four consecutive tests on the same system.

as guided by simulation and testing, and implement the computational steps in a way that minimizes processing time on hardware that is suitably fast.

A commonly used method for extraction of frequency information from a signal is the Fast Fourier Transform or FFT. While a powerful tool, they face drawbacks such as being computationally intensive, of particular relevance to this application. When computing the discrete Fourier Transform using a method such as the Cooley-Tukey FFT algorithm,[6] the relationship between the length of the input sample and the number of output bins generated is linear. Therefore, a requirement for adequately precise bin spacing would also represent a requirement on time spent collecting the sample, which may be in conflict with the need for minimal delay in a system's feedback response. Thus, an FFT-based approach will face challenges of finding a usable compromise between lag, ability to identify transients, and adequate frequency precision.

When the dynamics of a system are understood adequately, and the range in which responses will fall is known, some alternatives to the FFT can offer improved performance characteristics while avoiding some of the drawbacks identified. The following presents an investigation into a frequency measurement method that seeks to allow for tracking in high-rate systems, with comparison to an FFT-based method's frequency tracking performance.

One of the notable results is the proposal and demonstration of the Delayed Comparison Error Minimization frequency detection method. The theory behind the method is laid out, a functional implementation is developed, and the performance of that implementation is investigated besides, and compared to, the FFT approach. The results demonstrate that the method has applicability and some advantages in high-rate state estimation.

2. BACKGROUND

High rate dynamic events are characterized as occurring in less than 100 ms, along with possessing: 1) large uncertainties in the external load; 2) high levels of nonstationarities and heavy disturbances; and 3) generation of unmodeled dynamics from changes in system configuration.[1] These specific challenges are demonstrated and discussed by Hong et al. using representative experimental data.[7] This previously conducted experiment was performed by subjecting circuit boards and accelerometers assembled into an electronics package using

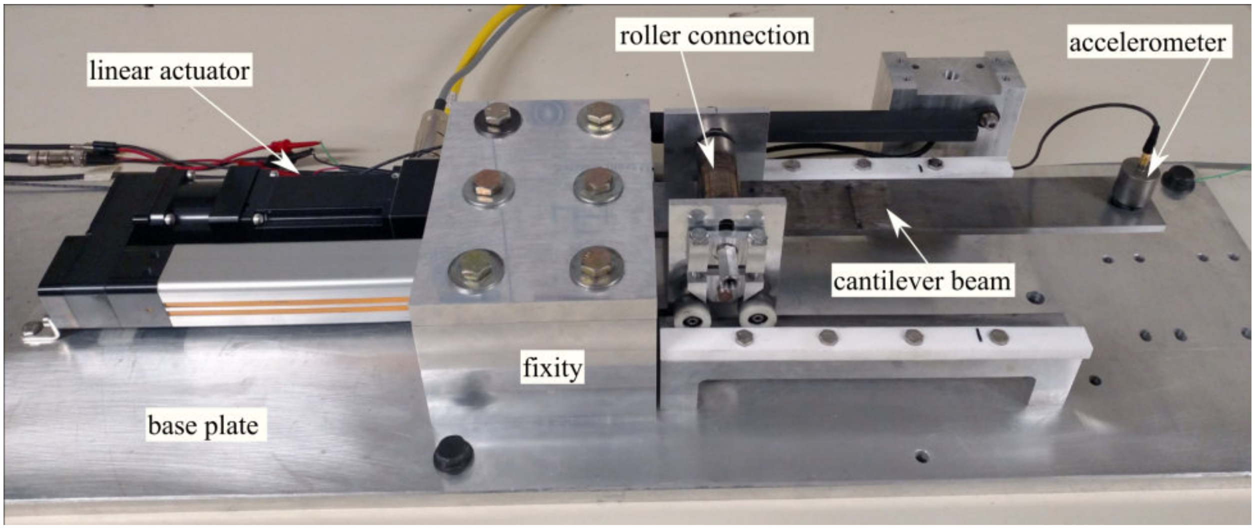

Figure 2. The Dynamic Reproduction of Projectiles in Ballistic Environments for Advanced Research (DROPBEAR) experimental testbed.

shock survivability methods to high impact conditions using an accelerated drop tower. The electronics test assembly underwent 4 consecutive assisted drops in the study. Some of the results are plotted in figure 1. The three challenges that characterize high-rate dynamic events are all shown in this study: 1) Unknown external load, 2) High level of nonstationarities (preventing meaningful determination of a running average) and heavy disturbances, and 3) Variations in the response of the structure even when performing tests sequentially. It demonstrates plainly the difficulties present in predicting high rate dynamic events, most notably due to the very short time scale on which the events occur, and the cumulative damage building between tests resulting in inconsistencies in response.

While some of the difficulties of dynamic high rate testing are visible when reviewing the previously discussed accelerated drop tower test series, the applicability to the development and experimental validation of observers able to perform rapid state estimation on structures experiencing high rate dynamic events is hindered by the inconsistencies resulting from test article damage. Recognizing this difficulty and seeking to address it, Joyce et al. introduced and modeled the Dynamic Reproduction of Projectiles in Ballistic Environments for Advanced Research (DROPBEAR),[8] which is seen in figure 2. The DROPBEAR testbed consists of a cantilever beam and a time varying parameter programmed by the user. The varying parameter is the position of a set of rollers on a linear actuator representing a continuously adjustable pin condition. The design intent is that the pin condition from the moving rollers represents a nonstationary boundary condition. While the input parameter seeks to represent changes such as damage occurring in a structural system, it is designed with the intent to produce controllable and repeatable changes in the results. The repeatability of this system was unobtainable in the previously performed accelerated drop tower tests, with the destructive nature of changes in the previous test vs the controlled changes on the DROPBEAR being responsible for this difference. The DROPBEAR was used by Downey et al. in the development of a millisecond model-updating technique, which updated the FEA model of the structural system through comparison of the model to the structural system in the frequency domain and adjusting to minimize the error between the two.[9] The results of this method showed that it was possible to update the pin location in the model every 4.04 ms with a 2.9% accuracy.[9]

The intent behind developing the DROPBEAR testbed is to serve as a generalized representation of various real systems in existence, with a non linear response to a continuously variable input parameter. The DROPBEAR's recording of actual, rather than only requested, location of the input parameter improves verification of the data generated on the testbed. Because of the advantages of the DROPBEAR system, its output was also used to guide generation of a simplified synthetic data set. As Downey et al. demonstrate in Ref.,[9] the acceleration data generated by the DROPBEAR's response to a dynamic input presents some challenges to being analyzed for frequency content using the FFT method described in more detail further in this paper, with a particular challenge in achieving adequate precision within the time allowed to not cause excessive lag. These simplified data sets seek to isolate the potentially challenging aspects of DROPBEAR data that make it difficult to analyze; then each of these aspects can be analyzed and a better understanding of the problem used to pursue solutions.

3. METHODOLOGY

Each of the subsections will explain how one of the frequency detection methods operates, characteristics of the output they generate, and how the function of each can be affected both by tuning parameters as well as certain characteristics of the data being fed to them.

Rolling FFT

One method for converting signals from the time domain to the frequency domain is the Fast Fourier Transform, or FFT for short. If the FFT algorithm is fed an even number of samples of real numbers, it will produce half that many frequency bins for the output. The bins will span from DC to half the sampling rate of the input, commonly called the Nyquist frequency, with even spacing between them. Therefore the sampling rate divided by the number of samples provided to the FFT algorithm will determine the spacing between adjacent frequency bins, which can be interpreted as the frequency precision of the output. The cumulative amplitude and duration of frequency components falling within a bin's upper and lower boundaries will determine the magnitude of the number that the bin contains. When considering only the magnitude of the frequency bins, the time at which an event occurred within a series of samples can't be determined. It is also worth noting that a long duration, low amplitude event can have a similar effect on FFT outputs as a very short, high amplitude event. These well-known limitations[9] are stated here again for emphasis of the difficulties that FFTs can face in trying to track the frequency of a signal that varies rapidly.

The FFT is being applied in this paper in order to determine the frequency response of a system with minimal delay. In order to determine when changes are occurring, and return this information rapidly, the FFT is repeatedly run on short windows that are passed over the time-series data. As the implementation being used bears similarity to shutter speed in photography, where an event happening at any time affects the output and, in the photography example, would cause blurring, a short window can help to determine when in time an event occurred. Analyzing shorter sections and running them more often can help to determine when in the time domain an event occurred, though the advantages will diminish as the windows overlap further and further. There is also the drawback that frequency precision will decrease with shorter window lengths. It becomes apparent that a compromise between temporal and frequency precision is necessary when attempting to use an FFT-based method. These tradeoffs are demonstrated in this paper by using a selection of FFT lengths to show how accuracy in the time domain and accuracy in the frequency domain are each affected, along with lag. For any that have considered the effect of increasing the sampling rate, it will increase the number of bins but that change is accompanied by the frequency range of the highest bin increasing proportionally; frequency range is increased, but bin spacing or frequency precision are not affected.

The primary frequency detection scheme in this paper works by discarding bins in the FFT output that are outside of the relevant/expected range and then picking the bin containing the highest absolute value to provide a simple interpretation of the frequency domain data. The timestamp for each fundamental frequency point is taken from the last input sample used by the FFT; this means that the timing is analogues to running the analysis in real-time, on theoretically perfect hardware that introduced no additional delay or processing time, representing only lag fundamentally inherent to the analysis technique.

Delayed Comparison Error Minimization

The Delayed Comparison Error Minimization method, like an FFT-based approach, identifies the primary frequency of a sample set. The method, originally proposed by Scheppegrell et al. in ref,[10] seeks to achieve similar or better frequency precision when working with fewer samples of time series data. This should allow the Delayed Comparison Error Minimization technique to generate a more precise and accurate estimate of the signal's primary frequency with less lag than an FFT-based frequency estimation system. The Delayed Comparison technique works by comparing sections of a periodic signal with a known time difference between them. If a signal has a periodic component, finding the time difference which results in the smallest difference between samples will indicate the signal's period. A similar technique can be seen in.[11]

This paper expands the previously developed Delayed Comparison Error Minimization method to consider signals with varying amplitude as well as varying frequency components. The specific implementation used in this paper normalizes the amplitude of the values in the two lists to be compared, compares each respective point to produce a difference list the same length as each input sample set, the values in the list are squared and then summed to a single value. This implementation of the concept gives a bandwidth extending from the frequency of a period with length equal to the longest time difference, or delay, between compared sections to half the frequency at which the data was sampled. Frequency precision is the inverse of the sampling rate of the provided sample data.

Normalization of the values in each list is performed by finding the max value contained in them, and then dividing all values in a list by a value proportional to the max. The normalization step is intended to improve matches between the lists when the amplitude of the input signal is changing with time, which ideally a primary frequency detection method would be unaffected by. Without normalization, the difference between the sample and comparison lists may be large even when the time difference between them is equal to the wave period, due to the difference in amplitude of each list; this may result in a skewed determination of the signal's primary frequency.

Algorithm 1 Pseudocode for the Delayed Comparison Error Minimization Method

```
collect 400 data points from Signal
for 300 Cycles do
    Reference = 100 most recent data points, all divided by amplitude of max value in list
    Comparison = 100 data points, starting from delay equal to the cycle count, all
divided by amplitude of max value in list
    Difference = point by point difference between Reference and Comparison
    DifferenceSquare = Difference*Difference
    ErrorSquareSum = sum of all points in ErrorSquare
    append ErrorSquareSum value to SumVsDelay
end for
SignalPeriod = position of minimum value between the 100th and 200th point in SumVsDelay
calculate Frequency using SignalPeriod and SamplingRate
```

In algorithm 1, pseudocode helping to explain the steps in the implementation of the Delayed Comparison Error Minimization concept is presented. Table 1 defines the variables used. Figures 3-6 provide a visual representation of the steps in the method used in order to ease understanding. The explanation of each step in the implementation follows.

1. Figure 3(a). A number of samples are chosen, beginning with the sample at Time = 0 in the data set (note that zero is on the right in these figures), through the sample number matching the sum of two parameters. The first parameter is the length of the longest period to be detected, the second the length of the "reference" and "comparison" sample sets used for the comparison. These values are, respectively, 300 and 100 samples, or 30 ms and 10 ms at the 10,000 samples/second sampling rate, in the implementation here.

2. Figure 3(a). The right-most 100 samples are copied into "reference", and on the first cycle, as in figure 3, those same 100 values are copied to "comparison".

3. Figure 3(b). The point-by-point differences between "reference" and the respective values in "comparison" are found, with the output being a list of 100 difference values

4. Figure 3(c). The square of each value in the list is calculated, then

5. All values in the list are summed, as seen in "current sum" in figure 3(d).

6. That value is stored as the first number in the "sum vs delay" list, containing all the difference squared sums versus delay values, and represents the difference at a delay value of 0.
7. Steps (2) through (6) are repeated, except that the samples placed in "comparison" are copied from one sample earlier, or to the left, with each subsequent pass. The sum values are each appended to the list, labeled "sum vs delay" in figure 3(d), with each additional number resulting from a delay 1 sample space greater. This is repeated until the last cycle, seen in figure 6, as the data points to populate "comparison" starts 300 points away from the beginning of the points going into "reference".

After the last cycle is completed and the "sum vs delay" list is populated, the local minimum within the expected range of wave periods is found. The position of the minimum indicates the period of the signal and is then copied to another list "signal period", along with the timestamp of the last data point used in that comparison cycle. The frequencies can readily be calculated from the signal periods. If a longer signal is to be analyzed, each step described up to this point is repeated at different sections of the signal until the complete signal is analyzed at the desired temporal resolution.

Table 1. Definitions for the components used in the Delayed Comparison Error Minimization technique.

Parameter	Definition
SamplingRate	Sampling rate used to collect data.
Signal	Periodic signal of unknown frequency.
Reference	List of the last 100 samples from Signal.
Comparison	List of 100 delayed samples from Signal.
Difference	List of Reference - Comparison.
DifferenceSquare	Error^2
DifferenceSquareSum	A single value representing the sum of values in the list ErrorSquare.
SumVsDelay	ErrorSquareSum values arranged by delay.
SignalPeriod	Signal period length, in samples.
Frequency	Frequency of Signal in Hz.

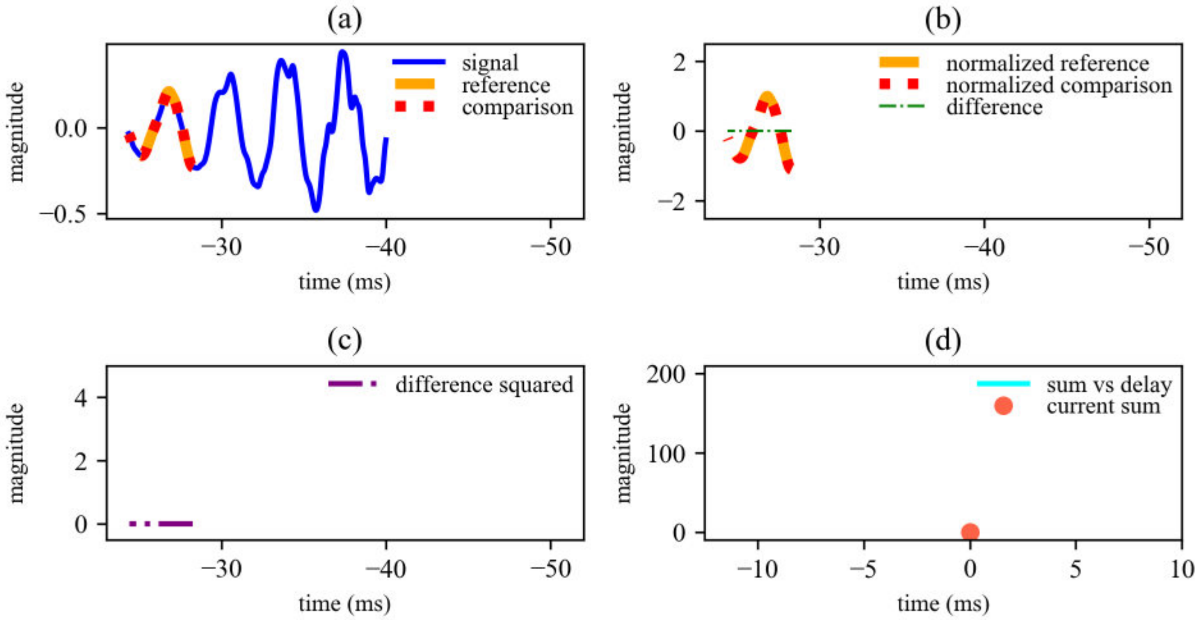

Figure 3. Demonstration of the steps to determine wave period; error evaluated at 0 ms delay. (a): signal, reference, and comparison (b): reference and comparison overlapped, difference (c): difference squared (d): sum vs delay, current sum

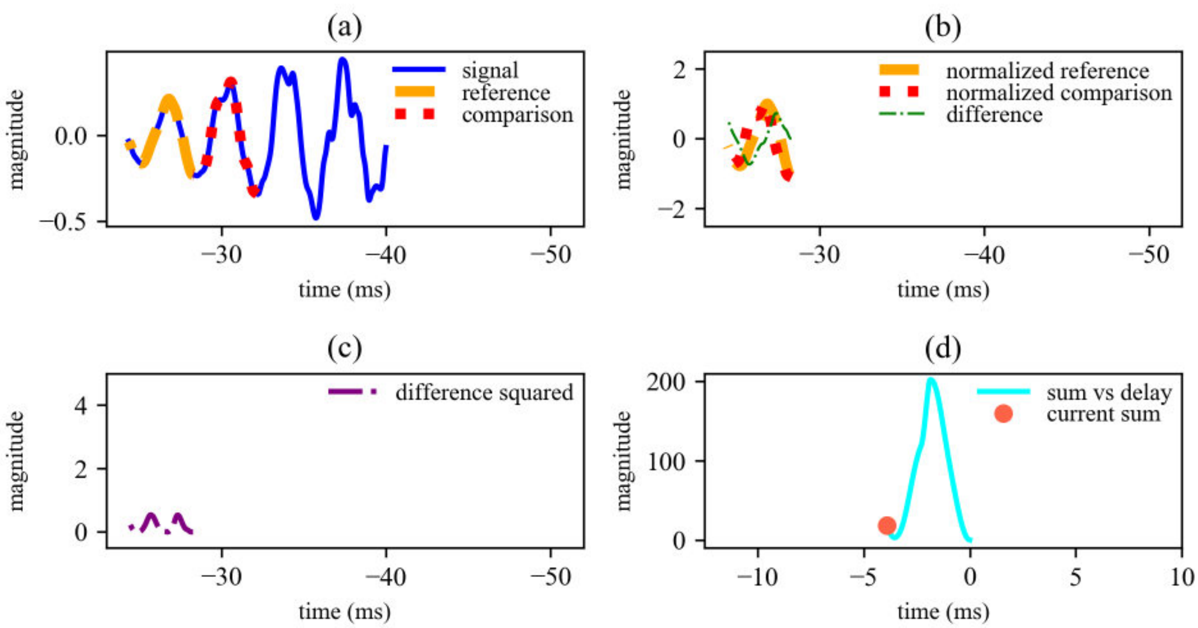

Figure 4. Demonstration of the steps to determine wave period; error evaluated at 10 ms delay. (a): signal, reference, and comparison (b): reference and comparison overlapped, difference (c): difference squared (d): sum vs delay, current sum

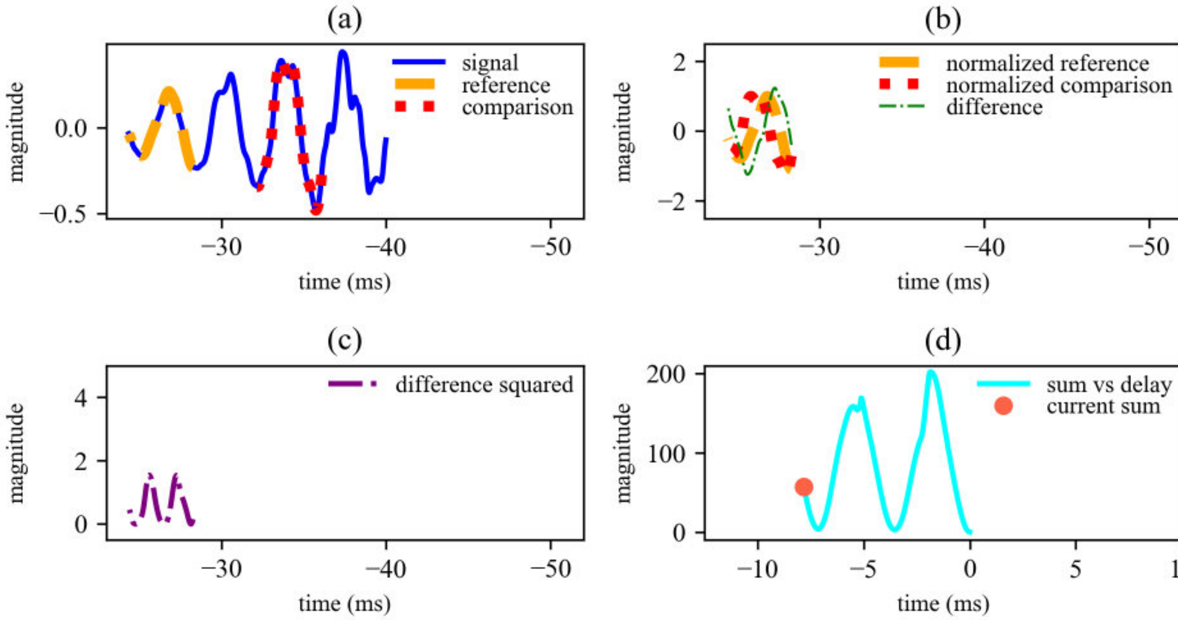

Figure 5. Demonstration of the steps to determine wave period; error evaluated at 20 ms delay. (a): signal, reference, and comparison (b): reference and comparison overlapped, difference (c): difference squared (d): sum vs delay, current sum

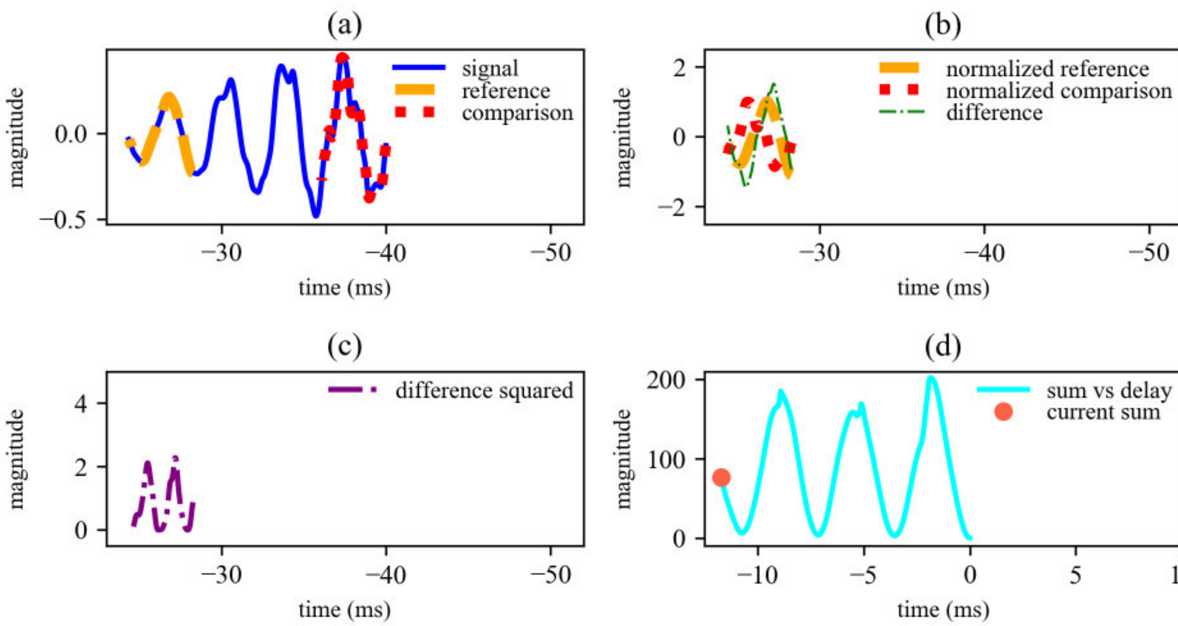

Figure 6. Demonstration of the steps to determine wave period; error evaluated at 30 ms delay. (a): signal, reference, and comparison (b): reference and comparison overlapped, difference (c): difference squared (d): sum vs delay, current sum

4. RESULTS

The performance of the two techniques was compared by providing them each with the same data sets and using common metrics to quantify their performance. First, a numerical analysis for a synthetic signal with a constant amplitude linear frequency sweep is performed. Next, a synthetic signal with a decreasing amplitude during a linear frequency sweep is investigated. Lastly, an experimental dataset collected from an accelerometer mounted to the DROPBEAR physical system as previously presented by Downey et al in ref[9] is considered. Metrics were defined to quantify the difference between the FFT-based method and the Delayed Comparison Error Minimization technique. The time difference between the most recent sample and the oldest sample used in calculations, will be calculated and referred to as the maximum theoretical delay. Observed delay, the subjective visual amount that the estimate lags behind the original signal, will also be noted and used for comparison between the methods and parameters.

Numerical Analysis

(a)

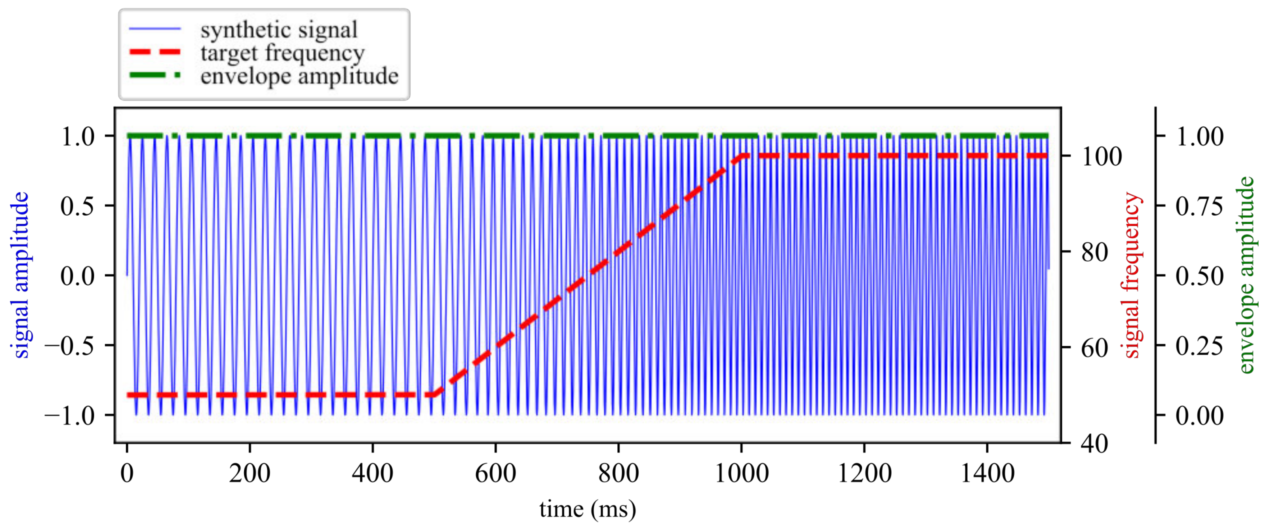

(b)

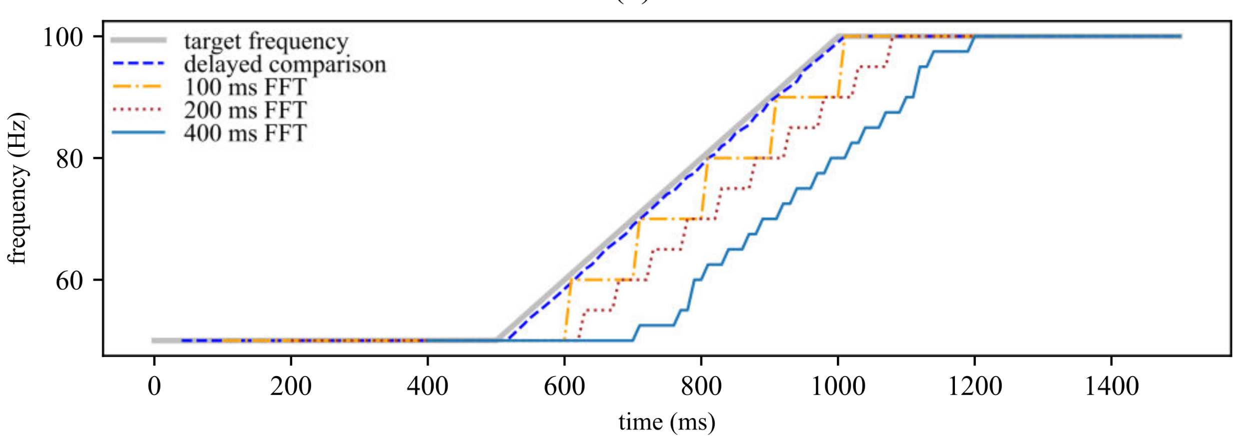

Figure 7. (a) Phase coherent constant amplitude sweep for analysis. (b) Plot demonstrating tracking of FFT and delayed comparison methods on constant amplitude sweep.

The first data set, seen in figure 7 (a), is constructed from half a second at 50 hz, then it undergoes a linear sweep from 50 hz to 100 hz over a half-second while maintaining phase coherence with the previous section, then the last section is 100 hz for half a second. The sampling rate is 1000 samples/second. The set contains only a single frequency component at any given time. For this dataset, the time difference between the most recent sample and the oldest sample used in calculations will be calculated and referred to as the maximum theoretical delay. Observed delay, the subjective visual amount that the estimate lags behind the original signal, will also be noted and used for comparison between the methods and parameters.

Figure 7 (b) shows the output of the Delayed Comparison and FFT-based methods using several window lengths. FFT lengths of 100 ms, 200 ms, and 400 ms result in frequency precisions of 0 hz, 5 hz, and 2.5 hz, respectively. Each point for the FFT traces are made using the samples from between the time they are plotted and the time earlier by the length of the FFT. The "steps" visually apparent in the traces are a result of the frequency precision limitations, according to how many frequency bins each FFT generates. With a longer FFT, the precision and number of bins in the FFT will increase, and the size of the steps will decrease. A longer FFT requires more samples, however, and results in more lag; this can be seen in the horizontal space between the target and estimated frequency plots. The FFT window length is also the age of the oldest sample used, thus the maximum theoretical lag is also equal to the length of the FFT. A subjective visual determination of the time difference suggested, at least to this individual, that observed lag is approximately half of the window length on the data set used here; amplitude is constant and frequency varies at a linear rate, results will likely vary when using data that doesn't share those characteristics.

As the synthetic dataset was known not to contain any frequency content below 50 hz, the tracker was set to minimize lag while covering that frequency on the low end. In this setup, the age of the oldest sample was the sum of the period of the lowest frequency (20 ms for 50 hz), plus a margin of 10 ms, plus the length of the samples compared (10 ms, when the samples are half the length of the lowest expected frequency's period). Adding these up gives 40 ms as the maximum theoretical delay. The visual subjective determination of lag by the author is around 10 ms; again, it is likely that a different dataset will affect that number. While the results contain some jitter that's not presently well understood, the lag, precision, and accuracy of this method all appear to be significantly better than is achieved by an FFT-based system at any of the length settings.

While the results above demonstrate the feasibility of the delayed comparison error minimization technique on data of a constant amplitude, it can be observed that the data collected from the DROPBEAR system varies in amplitude. As a frequency tracking method would ideally be able to operate on data collected from a system with varying outputs, sensitivity to amplitude changes should be tested for and mitigated or eliminated if possible within the expected range. In order to evaluate the error in frequency determination that results from amplitude variations of the input data, a synthetic sweep similar to the previous one was generated, but with an amplitude change introduced.

(a)

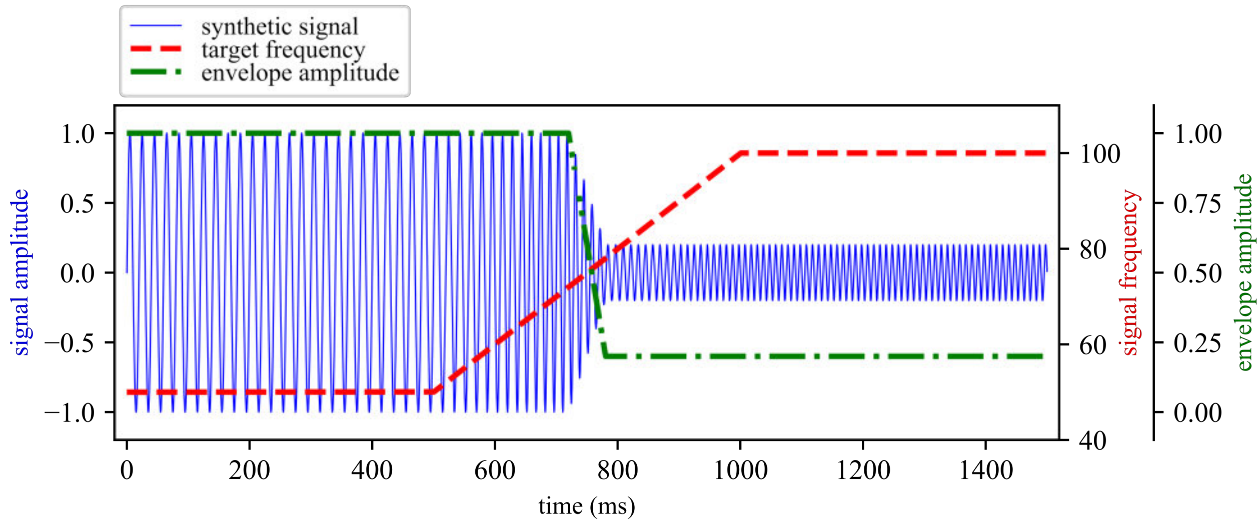

(b)

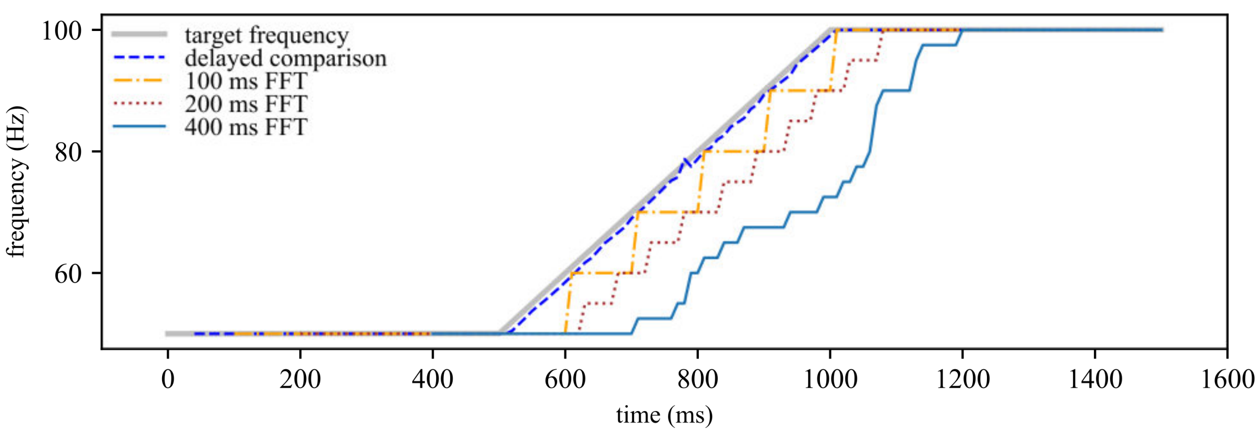

Figure 8. (a) Decreasing amplitude sweep for analysis. (b) Plot demonstrating tracking of FFT and delayed comparison methods on decreasing amplitude sweep.

Like figure 7(a), the sweep starts with a half-second of 50 hz, followed by a half-second transition from 50 hz to 100 hz, and ends with a half-second at 100 hz. The next sweep, shown in figure 8(a), begins at an amplitude of 1 and maintains that through 720 ms. From that time through 780 ms, the amplitude decreases linearly to 0.2, which it maintains through the end of the sweep.

When each of the previously described methods for frequency determination are applied, the result is as shown in figure 8(b). It can be seen that the Delayed Comparison method has developed an unusual looking, and inaccurate, "spike" around the 780 ms time, suggesting that the method struggles most as the sweep changes from a decreasing amplitude to a constant amplitude. Interestingly, the long FFT with a 400 ms window also appears significantly affected by the amplitude variation.

Experimental Analysis

Experimental data collected from an accelerometer on DROPBEAR system is also used to demonstrate the performance of each method. The plot shown in figure 9 displays an excerpt from an accelerometer mounted towards the end of the DROPBEAR's cantilever beam. The acceleration data collected from the end of the beam is overlaid with the position of the pin; frequency varies somewhat predictably with pin location within the displacement range used. At around 12500 ms to 13500 ms on 9, the shape of the wave envelope has a visible beat frequency. This suggests that there are at least two frequency components of similar frequency and amplitude within the signal at that point in time, presenting a possible challenge when determining the primary frequency of the signal.

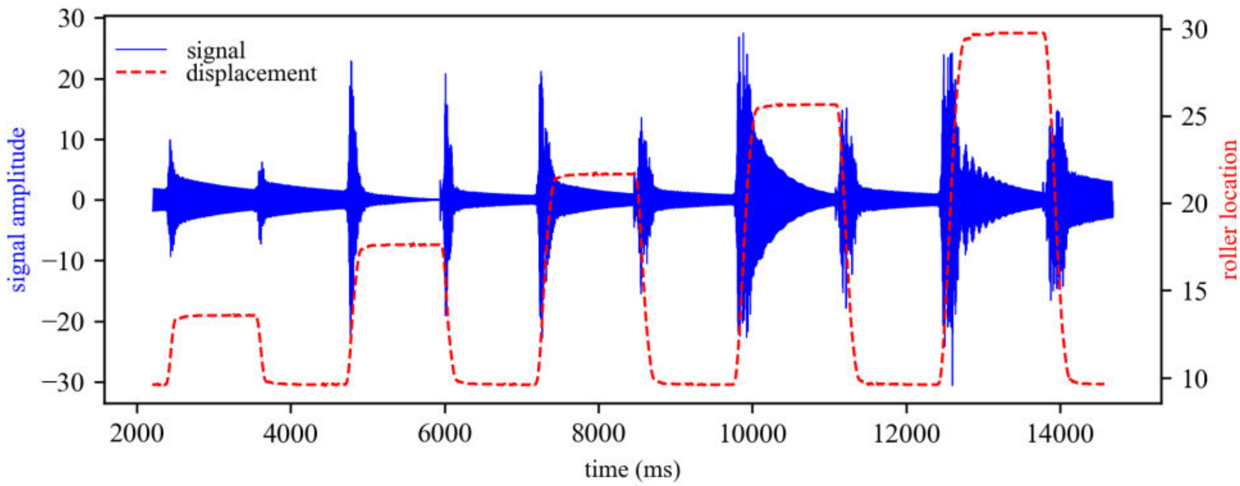

Figure 9. DROPBEAR roller location vs measured acceleration

Figure 10 shows the relative performance of the Delayed Compariso Error Minimization method and each FFT window length when running on data coming from a physical system. The Delayed Comparison Error Minimization method appears to pick up high-frequency noise and displays large output swings at several points in the trace, especially when the pin is moving; inspecting the signal at these times shows that high-frequency content is generated by the DROPBEAR system during pin motion, and it appears that the Delayed Comparison Error Minimization method's short samples result in it picking up on these much more visibly than the FFT methods. On the section of data that displays an apparent beat frequency, the Delayed Comparison Error Minimization method displays some hunting behavior where it switches back and forth between two similar frequencies. The Delayed Comparison Error Minimization method also displays some jumps in its output when the pin is stationary at a position that results in higher frequency oscillation; this is likely due to the Delayed Comparison Error Minimization method's difficulty in telling apart wave periods that are integer multiples of each other, further work will focus on determining if this drawback can be mitigated or solved.

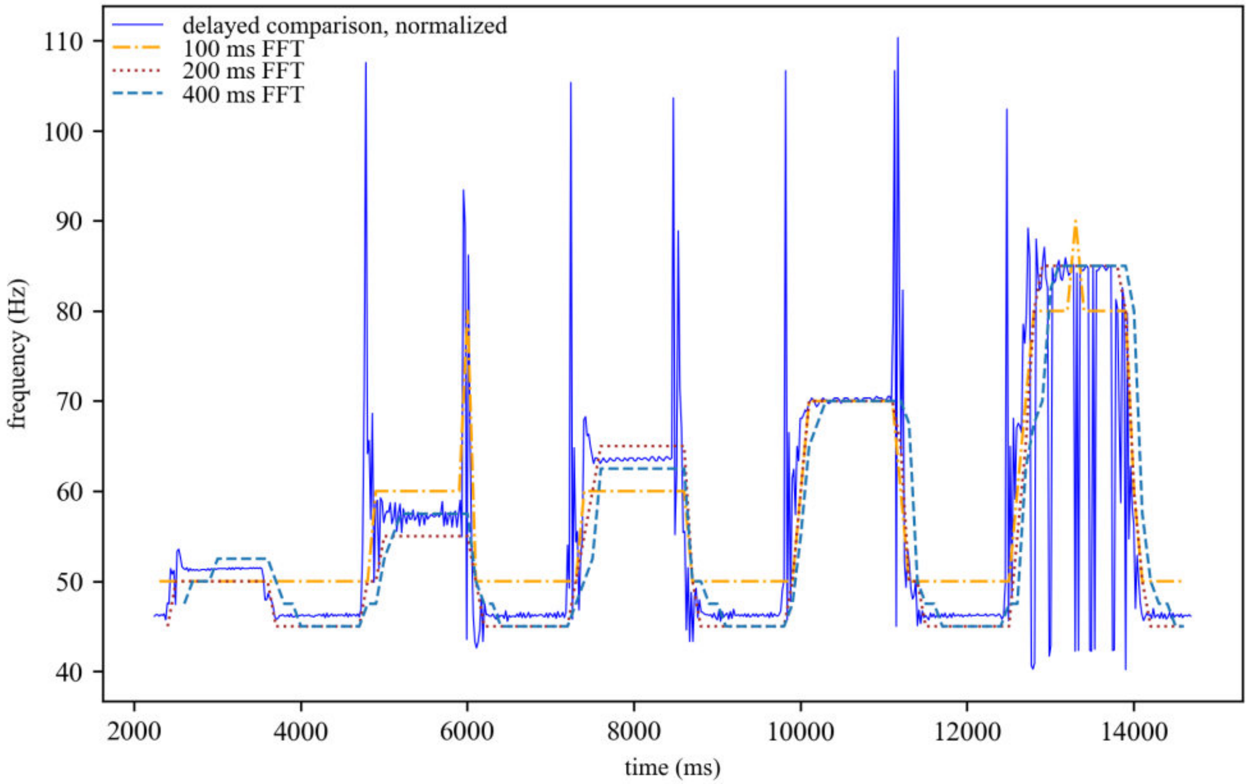

Figure 10. Delayed Error Comparison and FFT tracking on DROPBEAR system.

5. CONCLUSION

The performance of the Delayed Comparison Error Minimization method in the numerical analysis demonstrates its advantages in terms of both frequency precision and low lag. In the experimental analysis, the dataset used demonstrates that the current implementation of Delayed Comparison Error Minimization may not be the best choice in some applications.

A characteristic which may be interpreted as a disadvantage in some circumstances is that it sometimes reports frequencies which were present only briefly in the signal, while other techniques with longer sampling windows may be less likely to treat them as the dominant frequency. Whether this is considered an advantage or disadvantage may depend on the application, and this characteristic may be mitigated through variation of window length and filtering of the output, likely at the expense of increased lag thus diminishing the advantages of the technique.

Knowledge of the expected boundaries of the signal being analyzed is required in order to achieve maximumum performance, and there appears to be difficulty with distinguishing harmonics in the current implementation. If the delay parameter is set too short, the method will be unable to correctly detect frequencies below that period length. Improvement in distinguishing harmonics is expected through analysis and processing of the sum vs delay array.

Delayed Comparison Error Minimization is a tool that currently requires the signal on which it is operating to be somewhat known beforehand, such that its parameters can be tuned accordingly in order to operate correctly and reach maximum performance, which presently operates best on signals varying by less than an octave, and where the signal either lacks high amplitude short duration noise or its detection is desirable. Within these limitation, it is an interesting option which exhibits compelling performance in analyzing signals which are rapidly changing and require low-lag characterization.

acknowledgment

This material is based upon work supported by the United States Air Force through the AFRL/RWK contract number FA8651-17-D-0002 and FA8651-17-F-1026. This work is also partly supported by the National Science Foundation Grant numbers. 1850012 and 1937535. The support of these agencies is gratefully acknowledged. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the United States Air Force. (Distribution A. Approved for public release; distribution unlimited (96TW-2020-0118)).

REFERENCES

[1] Hong, J., Laflamme, S., Dodson, J., and Joyce, B., "Introduction to state estimation of high-rate system dynamics," *Sensors* **18**, 217 (jan 2018).

[2] Wadley, H., Dharmasena, K., He, M., McMeeking, R., Evans, A., Bui-Thanh, T., and Radovitzky, R., "An active concept for limiting injuries caused by air blasts," *International Journal of Impact Engineering* **37**, 317–323 (mar 2010).

[3] Dodson, J. C., Lowe, R. D., Foley, J. R., Mougeotte, C., Geissler, D., and Cordes, J., "Dynamics of interfaces with static initial loading," in [*Dynamic Behavior of Materials, Volume 1*], 37–50, Springer International Publishing (aug 2013).

[4] Stein, C., Roybal, R., Tlomak, P., and Wilson, W., "A review of hypervelocity debris testing at the air force research laboratory," *Space Debris* **2**(4), 331–356 (2000).

[5] Chen, W. and Hao, H., "Numerical study of blast-resistant sandwich panels with rotational friction dampers," *International Journal of Structural Stability and Dynamics* **13**, 1350014 (jul 2013).

[6] Cooley, J. W. and Tukey, J. W., "An algorithm for the machine calculation of complex fourier series," *Math. Comput.* **19**, 297–301 (1965).

[7] Hong, J., Laflamme, S., Cao, L., Dodson, J., and Joyce, B., "Variable input observer for nonstationary high-rate dynamic systems," *Neural Computing and Applications* (Dec. 2018).

[8] Joyce, B., Dodson, J., Laflamme, S., and Hong, J., "An experimental test bed for developing high-rate structural health monitoring methods," *Shock and Vibration* **2018**, 1–10 (jun 2018).

[9] Downey, A., Hong, J., Dodson, J., Carroll, M., and Scheppegrell, J., "Millisecond model updating for structures experiencing unmodeled high-rate dynamic events," *Mechanical Systems and Signal Processing* **138**, 106551 (April 2020).

[10] Scheppegrell, J., Moura, A. G., Dodson, J., and Downey, A., "Optimization of rapid state estimation in structures subjected to high-rate boundary change," in [*Proceedings of the ASME 2020 Conference on Smart Materials, Adaptive Structures and Intelligent Systems (SMASIS 2020)*], SMASIS2020–2306, ASME (2020).

[11] Dorran, D., "Pitch/period tracking using autocorrelation." available at https://dadorran.wordpress.com/2014/09/24/pitchperiod-tracking-using-autocorrelation/ (Sept. 2014).