Hardware Implementation of Nonstationary Structural Dynamics Forecasting

Puja Chowdhury[a], Austin R.J. Downey[a, b], Jason D. Bakos[c], Simon Laflamme[d], and Chao Hu[e]

[a]Department of Mechanical Engineering, University of South Carolina Columbia, Columbia, SC, USA
[b]Department of Civil and Environmental Engineering, University of South Carolina Columbia, Columbia, SC, USA
[c]Department of Computer Science and Engineering, University of South Carolina Columbia, Columbia, SC, USA
[d]Department of Civil, Construction, and Environmental Engineering, Iowa State University, Ames, IA, USA
[e]Department of Mechanical Engineering, University of Connecticut, Storrs, CT, USA

ABSTRACT

High-rate time series forecasting has applications in the domain of high-rate structural health monitoring and control. Hypersonic vehicles and space infrastructure are examples of structural systems that would benefit from time series forecasting on temporal data, including oscillations of control surfaces or structural response to an impact. This paper reports on the development of a software-hardware methodology for the deterministic and low-latency time series forecasting of structural vibrations. The proposed methodology is a software-hardware co-design of a fast Fourier transform (FFT) approach to time series forecasting. The FFT-based technique is implemented in a variable-length sequence configuration. The data is first de-trended, after which the time series data is translated to the frequency domain, and frequency, amplitude, and phase measurements are acquired. Next, a subset of frequency components is collected, translated back to the time domain, recombined, and the data's trend is recovered. Finally, the recombined signals are propagated into the future to the chosen forecasting horizon. The developed methodology achieves fully deterministic timing by being implemented on a Field Programmable Gate Array (FPGA). The developed methodology is experimentally validated on a Kintex-7 70T FPGA using structural vibration data obtained from a test structure with varying levels of nonstationarities. Results demonstrate that the system is capable of forecasting time series data 1 millisecond into the future. Four data acquisition sampling rates from 128 to 25600 S/s are investigated and compared. Results show that for the current hardware (Kintex-7 70T), only data sampled at 512 S/s is viable for real-time time series forecasting with a total system latency of 39.05 µs in restoring signal. In totality, this research showed that for the considered FFT-based time series algorithm the fine-tuning of hyperparameters for a specific sampling rate means that the usefulness of the algorithm is limited to a signal that does not shift considerably from the frequency information of the original signal. FPGA resource utilization, timing constraints of various aspects of the methodology, and the algorithm accuracy and limitations concerning different data are discussed.

Keywords: real-time, forecasting, structural dynamics, FPGA, time series, FFT

Further author information: (Send correspondence to Austin R.J. Downey)

Austin R.J. Downey : E-mail: austindowney@sc.edu

Active and Passive Smart Structures and Integrated Systems XVII,
edited by Jinkyu Yang, Proc. of SPIE Vol. 12483, 1248316
 · 0277-786X · doi: 10.1117/12.2658036

1. INTRODUCTION

Real-time model-based control of active structures operating in high-rate environments requires real-time time series structural response forecasting. For example, hypersonic, space, and military systems require active control within the microsecond (µs) timescale as dictated by the dynamics of the system.[1] Real-time model-based control of these structures would enable real-time decision-making that would boost structure survivability in these extreme environments by modifying mission goals and outputs to changing conditions. According to Hong et al.,[2] the high-rate problem is marked by:

1. significant external load uncertainty;
2. high levels of nonstationarities and heavy disturbances; and
3. produced dynamics from modifications to the system design.

In general, structures that experience high-rate dynamics have acceleration amplitudes higher than 100 g_n for a duration of under 100 ms.

Time series forecasting with high-rate dynamics is difficult as any approach used must be robust enough to operate with noisy sensor data.[3] Time series forecasting is accomplished by examining patterns in a variable (or the connections between variables) or developing a model and using either the learned pattern or model to forecast signals into the future. To analyze time series data, it is common practice to divide the monitored variable into the three categories of trend, nonstationary, and residual.[4] Various methods can be used to do this, including sliding window, smoothing, and autoregressive expectation, which is widely used in forecasting high-rate dynamic system states, financial turn of events, environmental change, and energy interest.[5] When numerous time series components are present and their interactions need to be taken into account, a multivariate time arrangement is used.[6]

The timing requirements driven by µs structural health monitoring were articulated by Dodson et al.[1] Based on the dynamics of the considered "high-rate" class of systems, this work sets a system latency and forecasting horizon of 1 ms. To expand, the algorithmic work developed in this paper seeks to forecast the dynamic structural response (i.e. signal) 1 ms into the future while completing all required computations within a latency of 1 ms. To enable deterministic and low-latency time series forecasting of nonstationary signals, an FFT-based forecasting approach was developed that was implemented on a Field Programmable Gate Array (FPGA).

This study outlines the development of an online structural vibration time series forecasting hardware/software system.[7] The FFT-based forecasting technique is employed in this study and is implemented in a variable-length sequence configuration. After the data has been de-trended, it is translated into the frequency domain and measurements of frequency, amplitude, and phase are made. The data's trend is then retrieved by gathering a selection of frequency components, translating them back into the time domain, and recombining them. The

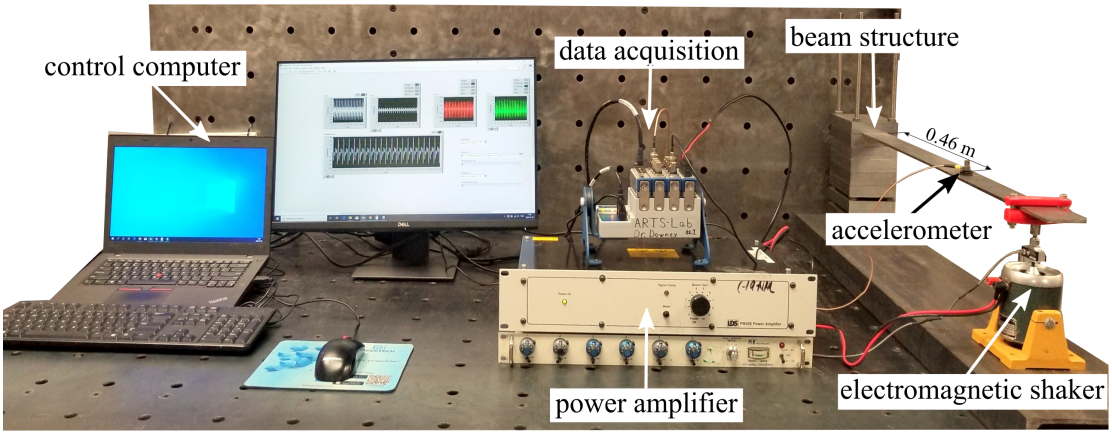

Figure 1. Setup for an experimental cantilever beam including main components and data acquisition.

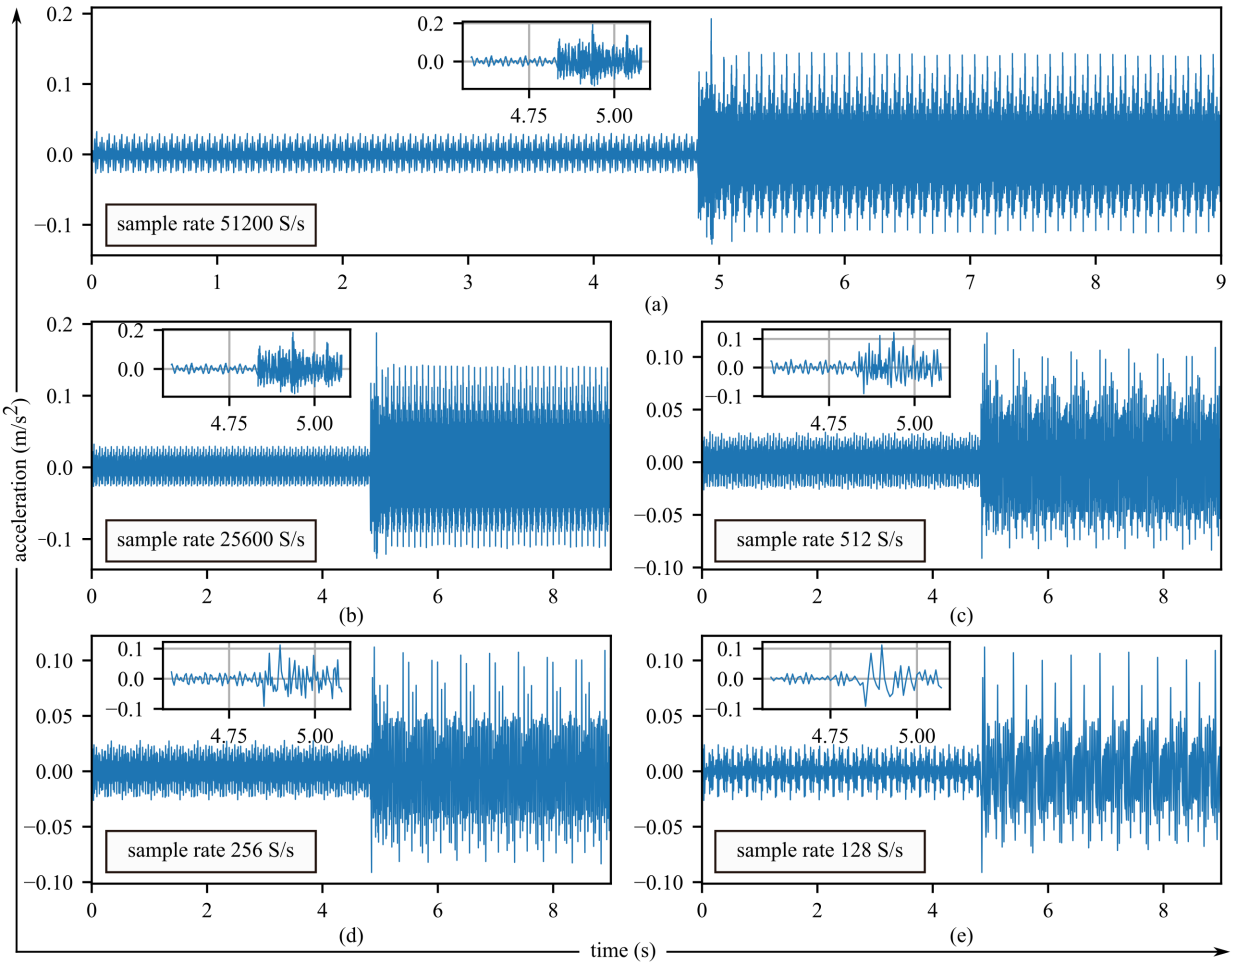

Figure 2. Data set with varied sample rates, showing: (a) the native sample rate of 51200 S/s; (b) sub-sampled at 25600 S/s; (c) 512 S/s; (d) 256 S/s; (e) 128 S/s; and inset plots provide a close look around nonstationary for each sampled data.

signals are extended into the future to the selected forecasting horizon at this point. The proposed methodology is experimentally evaluated on a Kintex-7 70T FPGA using structural vibration data from a test structure with varying levels of non-stationaries.[8] Four data collection sampling rates, ranging from 128 to 25600 S/s, are examined and compared.

Results show that the system can forecast time series data within the 1 ms latency constraint. However, it needs to be noted that a key challenge of FFT-based time series forecasting is that the periodicity of the time series signal must be properly considered. There must be sufficient perceptions of a period arrangement that may need to be processed as whole components (not partial cycles). To expand, there are challenges when trying to forecast a period signal when considering anything other than full periods of the signal that initiate and terminate as the zero-crossing. The contributions of this work are two-fold, 1) An experimental investigation showing the potential of the FFT-based time series forecasting methodology for high-rate signals, and 2) a detailed discussion of the periodicity challenge for FFT-based time series forecasting.

2. METHODOLOGY

This section describes the experimental testbed, experimental data, and the formulation of the FFT-based forecasting methodology.

2.1 Experimental Testbed

Figure 1 displays the experimental configuration used to develop experimental data for this work. The beam is excited by an electromagnetic shaker (model V203R manufactured by LDS), with a useful frequency range of 5-13000Hz and a peak sine force of 17.8N, and is driven by a power amplifier (model PA25E-CE manufactured by LDS). A 45 N load cell (model MLP-10 manufactured by Transducer Techniques) is mounted in-between the shaker and beam structure. A 24-bit bridge input signal conditioner (NI-9237 manufactured by National Instruments) is used to acquire the load-cell data. The experiment is run through a control computer with a Virtual Instrument written in LabVIEW.

Figure 2 reports the structure's measured acceleration response (x_v) for a composite sinusoidal input from the shaker. In this work, the composite signal is made up of 50, 70, and 100 Hz sinusoidal signals. Two sine wave signals are concatenated together at t=5 s where a nonstationary is present due to a change of frequency. To achieve this, an input signal of 0.25 V is used before t=5 s while a signal of 0.25 V is used after t=5 s. The first half of the composite signal is built from 50, 70, and 100 Hz frequencies while the second half signal consists of 50 and 100 Hz frequencies. Four different sampled data were created from this data and Figure 2 shows all of that including zoomed section near by nonstationary event. The original sampling rate of the data is displayed in figure 2 (a). This data is available in a public repository[8]

2.2 Algorithm Formulation for FFT-based Forecasting

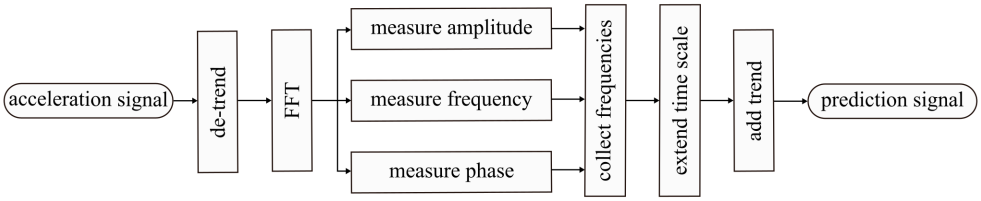

Figure 3. Schematic Algorithm diagram of FFT-based time series forecasting algorithms.

Figure 3 diagrams the algorithm used in this work for periodic structural vibration forecasting. The signal of observed acceleration is $x_{\mathrm{v}} = (x_1,\ x_2,\ x_3,\ \ldots\ , x_V)$ where V is the total sample points in the observed signal. A variable length sequence, x_{a} of size, N moves forward through time as time progresses. By applying the FFT-based time series forecasting method, a signal is generated that is M points long where $M > N$. The difference, $(M - N)$ presents the length of the forecasting horizon. By determining N and M, this method can be applied to achieve a predicted signal of desirable length. The variable length sequence is $x_{\mathrm{a}} = (x_{\mathrm{a}1},\ x_{\mathrm{a}2},\ x_{\mathrm{a}3},\ \ldots\ , x_{\mathrm{a}N})$. The first step is to remove any trend line from the acceleration x_{a}. To do this, a polynomial function is used where

$$x_{\mathrm{trend}} = p(x) = c_0 + c_1 x + c_2 x^2 + \cdots + c_q x^q \tag{1}$$

and q is the degree of the polynomial and c is a set of coefficients. In this work, $q = 1$. After removing the trend, the new acceleration signal without trend is $x = x_{\mathrm{a}} - x_{\mathrm{trend}}$ which has the same sample size as N. As considered, the acceleration signal without the trend, $x = (x_1,\ x_2,\ x_3,\ \ldots\ ,\ x_N)$, is a time series of N-samples that the frequency content is extracted from.[9] Therefore, the discrete Fourier transform (DFT) of that series can be expressed as

$$X_k = \sum_{n=0}^{N-1} x_n e^{(-i2\pi(kn/N))} \text{ for } k = 0,\ \ldots\ ,\ N \tag{2}$$

where,

$$\omega = 2\pi/N = 2\pi f \tag{3}$$

$$(X_{\text{amp}})_k = |X_k| \tag{4}$$

$$(X_{\text{phase}})_k = X_k/|X_k| \tag{5}$$

Similarly, the inverse DFT can be written as

$$x_n = \frac{1}{N}\sum_{k=0}^{N-1} X_k e^{(i2\pi kn/N)} \text{ for } n = 0, \ \ldots \ , \ N \tag{6}$$

Now, consider a new series of M length where $M > N$. Using amplitude and phase information, the time series can be constructed and written as

$$x_m = \sum_{k=0}^{M-1} ((X_{\text{amp}})_k \cos(2\pi(\dot{k}m/M)) + (X_{\text{phase}})_k) \text{ for } m = 1, \ \ldots \ , \ M \tag{7}$$

The x_m time series with the trend information added back can be expressed as

$$x_{\text{a\_new}} = x_m + x_{\text{trend}} \tag{8}$$

The FFT-based algorithm works best when the waveform is not interrupted. Let's say a signal is embedded with 1 Hz, 5 Hz, and 10 Hz frequencies. Now if an acquisition (learning) window of 1.5 seconds in length is considered, then all of these waveforms are cropped. In FFT, the time domain and frequency domain maintain the circular topologies. So, the two endpoints of input length are assumed to meet at the same point. But this is not true for this example. That's why it is necessary to ensure that the acquisition length considered for FFT must contain the integer number of periods. In a non-stationary signal, it is not possible to have all the embedded signals with different frequencies start at the same time. So, even taking a 1-second window cannot solve the situation. When the input length is shorter than the period of the lowest frequency component of the signal, periodicity in the predicted signal develops. However, when the input length exceeds the base period in the signal, the periodicity in the predicted signal is removed. Therefore, to capture this frequency, the minimum learning window length needs to be twice the period of the signal, per the Nyquist Theorem. For accurately capturing all the frequencies, the minimum period should be higher than the Nyquist limit.

2.3 Hardware Validation

While implementing any algorithm in FPGA hardware is a challenge, the FFT-based forecasting approach is a relatively simple algorithm making it well-suited for hardware implementation. In this work, hardware validation is done on a Kintex-7 70T FPGA housed in a NI cRIO-9035 that also incorporates a CPU running NI Linux Real-Time. Figure 4 diagrams how data is collected and processed on the FPGA, as well as how data is transmitted through parallel FFT-based forecasting. The sampling rate of the hardware system is set from 128 to 51,200 S/s and is restricted to intervals of the internal clock of the 24-bit ADC used in this project. Data is passed from the DAQ to FIFO and stored in the FPGA's look-up table memory. From FIFO through a for-loop data is going to

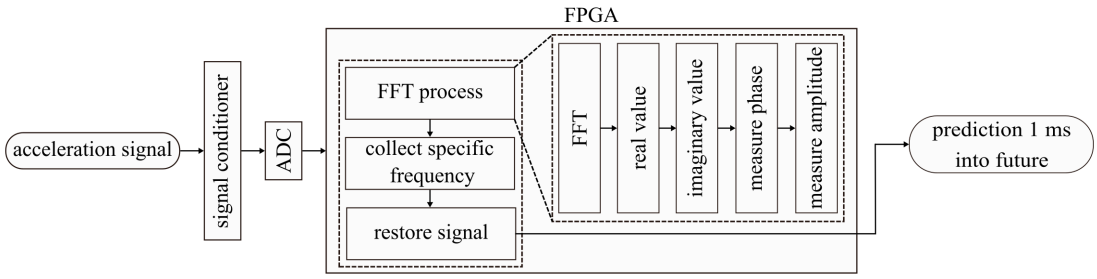

Figure 4. Flowchart for data collection and processing during FFT-based forecasting in case of hardware implementation.

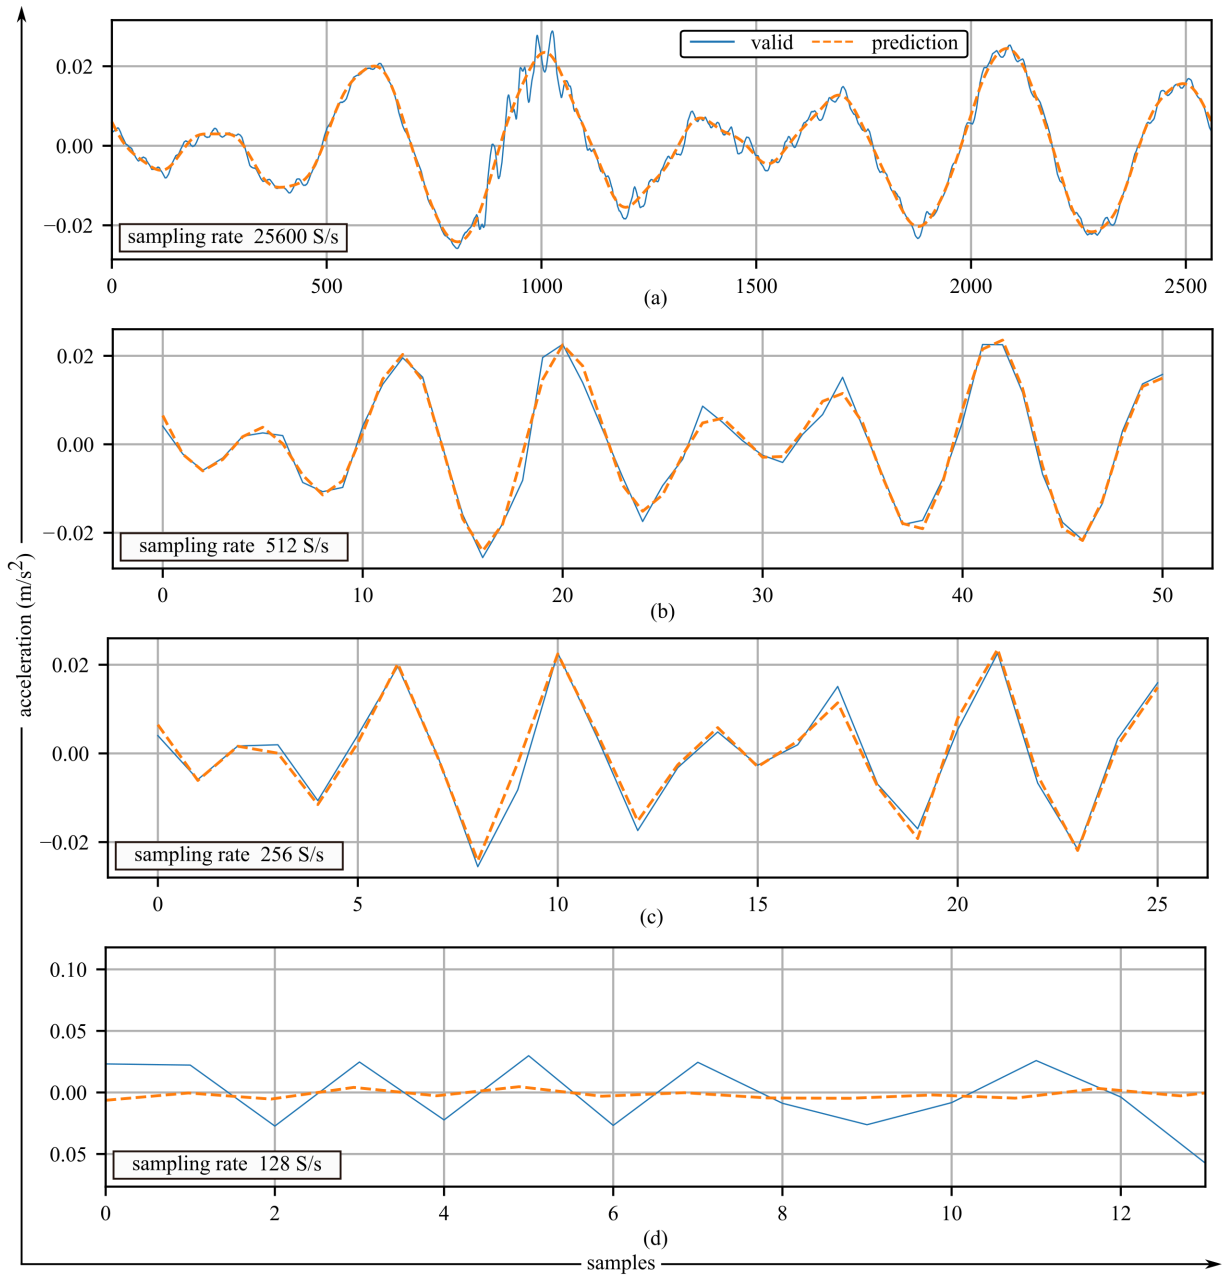

Figure 5. Simulation outcomes of forecasting at various sample rates, showing: (a) 25600 S/s; (b) 512 S/s; (c) 256 S/s; and (d) 128 S/s.

the FFT process. The next step is collecting specific frequencies to make more accurate forecasting and finally restoring the signal which is equivalent to 1 ms into the future is produced. The FFT process includes different steps like measuring real and imaginary, measuring phase, and measuring amplitude.

Table 1. FFT size and input length for different sampled data in hardware implementation.

sampling rate (S/s)	25600	512	256	128
FFT size	128	512	256	128
input (samples)	256	512	256	128

In this work, the LabVIEW FPGA development environment was used for developing the FPGA hardware designs, before being converted to a bitstream file through a Xilinx/Vivado workflow. The built-in LabVIEW FPGA FFT function has a range of size limitations between 8 to 8192 samples. Each size of FFT has a latency of cycles from 16 to 16384. For each sample rate, the goal was to pass a second of data to the FFT. However, due to hardware restrictions related to the chosen Kintex-7 70T FPGA, at the higher sampling rate of 25600 S/s the FPGA design could not meet timing requirements. Therefore, as shown in table 1, the FFT size and number of inputs are constrained for the sampling speeds of 25600 S/s. As a result, for subsequent research, only sampling data with a speed range of 128 to 25600 S/s is taken into consideration. The investigation of algorithm deployment on larger hardware is left to future work.

Table 2. For various sampling data, simulation outputs including foretasted signal RMSE, SNR, and chosen frequencies.

sampling rate (S/s)	RMSE	SNR	frequency list
25600	0.001727	17.12	50, 70, 100, 210, 220, 240, 260, 280, -50, -70, -100, -210,-220, -240, -260, -280
512	0.001889	16.33	50, 70, 100, -50, -70, -100
256	0.001911	16.18	50, 70, 100, -50, -70, -100
128	0.033853	0.15	50, 58, 22, 14, 20, 24, -50, -58, -22, -14, -20, -24

3. RESULTS

Figure 5 reports on the time series forecasting for four different sampling rates. Note that the native sampling rate of 51200 S/s is not shown for brevity as it performs similarly to 25600 S/s. Compared to the higher sampled data, the prediction accuracy for the lowest sampled data, 128 S/s, is poor. Note the significant drop-off in the algorithm's capabilities between 256 S/s and 128 S/s; which demonstrates that 256 S/s is the lower limit in terms of forecasting capabilities due to the loss of the higher frequency content in the signal down-sampled to 128 S/s.

Table 2 shows the frequencies used in reconstructing the signal and reports the RMSE and SNR for the four considered sampling speeds. The average difference between values predicted by a model and the actual values is measured by the Root Mean Squared Error (RMSE). The ratio of a signal's (valid) power to background noise (error) is known as the signal-to-noise ratio (SNR). Here, signals are expressed using the logarithmic decibel (dB) scale as the signals considered have a wide dynamic range. In contrast to other speeds such as 128 S/s, the frequency list reveals that 25600 S/s utilized more frequencies. When data is sampled at faster speeds, more frequency-rich information is gathered, allowing for frequencies to be used in the reconstruction of the signal. Due to this, the algorithm is considered useless at the lowest sampling speed of 128 S/s. Figure 6 reports the same RMSE and SNR data but in a graphical format, showing how values change in response to changes in the sampling speed.

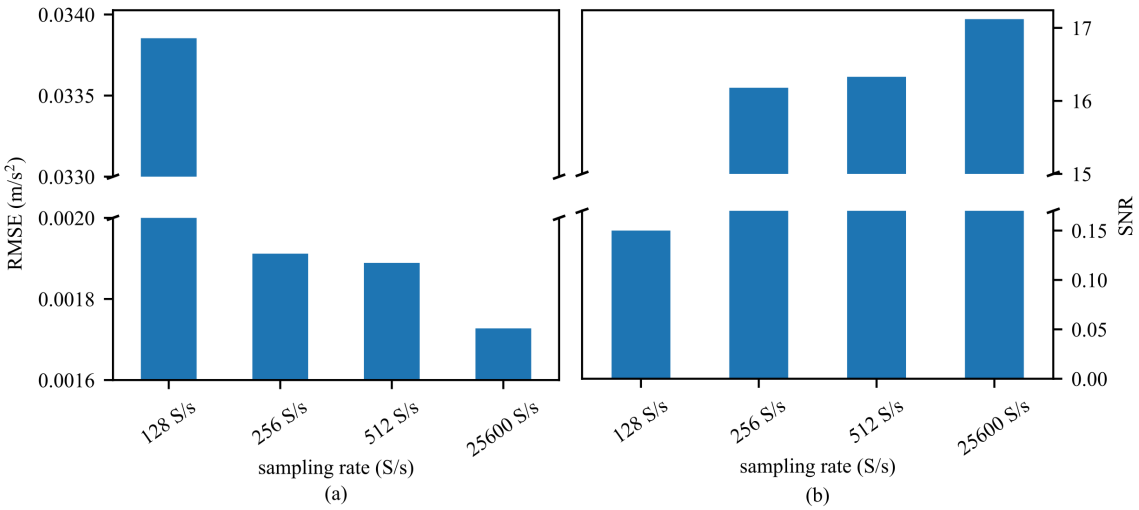

Figure 6. Effects of variously sampled data from simulation results, showing: (a) RMSE; and (b) SNR.

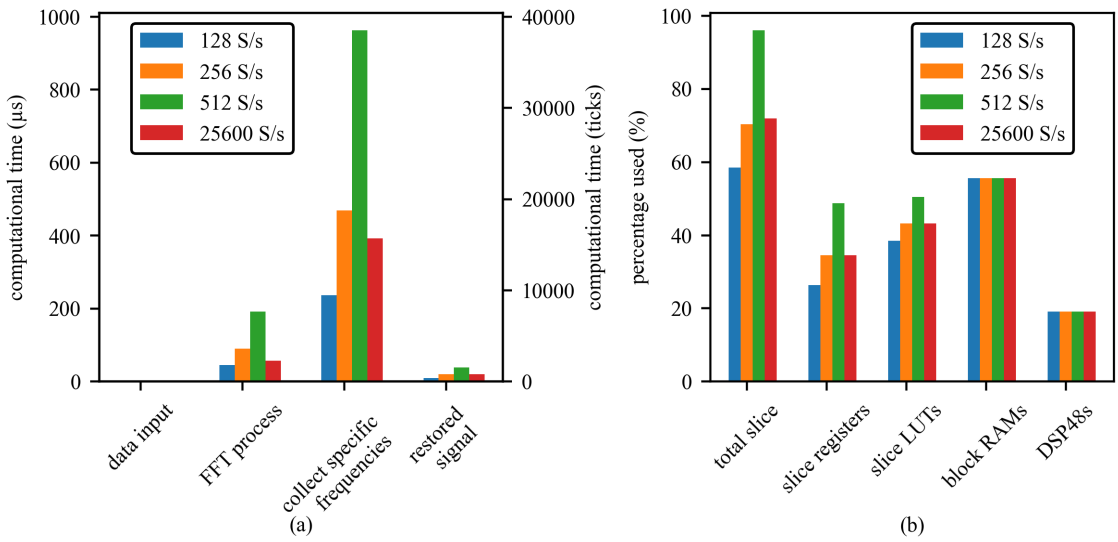

Figure 7. Results of the hardware validation procedure for varied sampled data in cases of (a) computation time required; and (b) device utilization.

Figure 7 shows the required computation time and device utilization for different intermediate steps of hardware implementation. The 512 S/s sampling rate takes greater computation time for various intermediate hardware implementation phases than other sampling rates, as seen in figure 7(a). Moreover, the larger the FFT size as defined in table 1, the larger the latency. This is the reason that the smallest sampled data 128 S/s has the lowest computation time, despite its forecasting results being completely unusable. Device utilization is shown in figure 7(b) and experiences the same situation where the faster sampling rates generally require more FPGA resources. Note that for device utilization, the signal sampled at 512 S/s uses 96% of the FPGA slices, signifying that a sample rate of 512 S/s, along with an FFT size of 512 samples (see table 1) is effectively the largest useful implementation of the FFT-based time series forecasting algorithm that can be deployed on the considered hardware (Kintex-7 70T).

Table 3. Time required for different aspects of FFT-based forecasting.

sampling rate (S/s)	input (samples)	data input		FFT process		collect specific frequency		restore signal	
		ticks	microsecond (µs)	ticks	microsecond (µs)	ticks	microsecond (µs)	ticks	microsecond (µs)
25600	256	1	0.025	2305	57.625	15717	392.925	802	20.05
512	512	1	0.025	7679	191.975	38502	962.55	1562	39.05
256	256	1	0.025	3584	89.6	18789	469.725	802	20.05
128	128	1	0.025	1792	44.8	9445	236.125	410	10.25

Table 3 illustrates that all data sampling speeds require the same amount of time for the data input step of 0.025 µs. An important outlier to note is that the sampling speed of 512 S/s requires the most time with a total latency of 39.05 µs in case of restoring signal. This is because the sample rate of 512 S/s is paired with an FFT size of 512; which maximizes the device hardware. In comparison, the higher sampling rates of 25600 required its pairing with reduced FFT sizes to enable its deployment on the chosen FPGA hardware.

Table 4. Device utilization for FFT-based forecasting where FPGA elements are shown by device utilization.

sampling rate (S/s)	total slice			slice registers			slice LUTs			block RAMs			DSP48s		
	slices used	slices available	**% used**	slices used	slices available	**% used**	slices used	slices available	**% used**	slices used	slices available	**% used**	slices used	slices available	**% used**
25600	7377	10250	72	28321	82000	34.5	17728	41000	43.2	75	135	55.6	46	240	19.2
512	9837		96	40052		48.8	20688		50.5				46		
256	7220		70.4	28320		34.5	17716		43.2				46		
128	5999		58.5	21663		26.4	15802		38.5				46		

Table 4 shows that the device utilization for DSP48s and block RAMs remains constant across all data. Except for this, every other FPGA component shows a considerable amount of variance depending on the sampling rate of the data.

4. CONCLUSION

This work describes the creation of a hardware/software system for real-time structural vibration time series forecasting. The suggested method makes use of a forecasting algorithm based on FFT. While any algorithm is difficult to implement in FPGA hardware, the straightforward nature of this technique makes it less complicated for hardware implementation. The FFT-based approach collects, processes, and extends the chosen frequencies to the forecast horizon. In FFT, the circular topologies are maintained in both the time domain and the frequency domain. Thus, it is assumed that the input length's two ends match. Because of this, it is essential to make sure that the acquisition duration taken into account for FFT must include an even number of periods that start and stop at zero. Results show that for the current hardware (Kintex-7 70T), only data sampled at 512 S/s is viable for real-time time series forecasting of the considered system with a total system latency of 39.05 µs in restoring signal. While a sampling speed of 256 S/s shows usefulness, a sampling speed of 25600 S/s requires FPGA resources beyond that provided by the chosen hardware. Lastly, the FFT-based time series algorithm itself completely falls apart for a sampling speed of 128 S/s. In totality, the tuning of hyperparameters for the FFT-based time series algorithm and its deployment onto FPGA hardware was found to be finicky and laborious while being tied to a signal within a limited frequency bandwidth. Future work will investigate the deployment of a hardware-in-a-loop implementation of the hardware/software system proposed here.

ACKNOWLEDGMENTS

The National Science Foundation provided support for this work through Grants 1850012 and 1937535. The National Science Foundation's support is sincerely thanked. The authors' opinions, results, conclusions, and recommendations in this material are their own and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] Dodson, J., Downey, A., Laflamme, S., Todd, M. D., Moura, A. G., Wang, Y., Mao, Z., Avitabile, P., and Blasch, E., "High-rate structural health monitoring and prognostics: An overview," in [*Data Science in Engineering, Volume 9*], 213–217, Springer International Publishing (oct 2021).

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

[3] Kay, S. M. and Marple, S. L., "Spectrum analysis—a modern perspective," *Proceedings of the IEEE* **69**(11), 1380–1419 (1981).

[4] Pang, Y.-H., Wang, H.-B., Zhao, J.-J., and Shang, D.-Y., "Analysis and prediction of hydraulic support load based on time series data modeling," *Geofluids* **2020** (2020).

[5] Owen, J., Eccles, B., Choo, B., and Woodings, M., "The application of auto–regressive time series modelling for the time–frequency analysis of civil engineering structures," *Engineering Structures* **23**(5), 521–536 (2001).

[6] Brockwell, P. J., Brockwell, P. J., Davis, R. A., and Davis, R. A., [*Introduction to time series and forecasting*], Springer (2016).

[7] Chowdhury, P., Barzegar, V., Satme, J., Downey, A. R., Laflamme, S., Bakos, J. D., and Hu, C., "Deterministic and low-latency time-series forecasting of nonstationary signals," in [*Active and Passive Smart Structures and Integrated Systems XVI*], **12043**, 466–472, SPIE (2022).

[8] Chowdhury, P., Downey, A., Bakos, J. D., and Conrad, P., "Dataset-4-univariate-signal-with-nonstationarity," (Apr. 2021).

[9] Chowdhury, P., Conrad, P., Bakos, J. D., and Downey, A., "Time series forecasting for structures subjected to nonstationary inputs," in [*Smart Materials, Adaptive Structures and Intelligent Systems*], **85499**, V001T03A008, American Society of Mechanical Engineers (2021).