

FEASIBILITY OF FRI-BASED SQUARE-WAVE RECONSTRUCTION WITH QUANTIZATION ERROR AND INTEGRATOR NOISE

Bryan He^{*‡} Alexander Wein^{†‡} Lakshminarayan Srinivasan[‡]

^{*} Department of Computer Science, California Institute of Technology

[†] Department of Mathematics, Massachusetts Institute of Technology

[‡] Neural Signal Processing Laboratory, Stanford University

ABSTRACT

Conventional Nyquist sampling and reconstruction of square waves at a finite rate will always result in aliasing because square waves are not band limited. Based on methods for signals with finite rate of innovation (FRI), generalized Analog Thresholding (gAT- n) is able to sample square waves at a much lower rate under ideal conditions. The target application is efficient, real-time, implantable neurotechnology that extracts spiking neural signals from the brain. This paper studies the effect of integrator noise and quantization error on the accuracy of reconstructed square waves. We explore realistic values for integrator noise and input signal amplitude, using specifications from the Texas Instruments IVC102 integrator chip as a first-pass example because of its readily-available data sheet. ADC resolution is varied from 1 to 16 bits. This analysis indicates that gAT-1 is robust against these hardware non-idealities where gAT-2 degrades less gracefully, which makes gAT-1 a prime target for hardware implementation in a custom integrated circuit.

Index Terms—square wave, compressed sensing, sub-Nyquist, finite rate of innovation, neurotechnology, analog thresholding, massive-scale neural recording, electrode arrays, optical methods

1. INTRODUCTION

The Nyquist-Shannon sampling theorem states that a signal that contains frequency components less than f Hz can be completely reconstructed from a sequence of measurements spaced less than $\frac{1}{2f}$ seconds apart. Direct sampling of the signal at less than $2f$ Hz results in distortion in the reconstructed signal, called aliasing. In the case of square waves, which contain non-zero amplitudes at infinitely large frequencies, no finite sampling frequency allows perfect reconstruction of the original signal from its samples using standard Nyquist sampling and reconstruction.

However, alternative methods of sampling and reconstruction based on compressive sensing [1] and signals with finite rate of innovation [2] do allow square waves to be sampled at much lower frequencies without aliasing. Signals with a finite rate of innovation are described by a finite number of parameters per unit time. This simplifying constraint allows various iterative techniques to recover the original signals from samples obtained at lower-than-Nyquist rates. Techniques for sub-Nyquist sampling for the related problem of level-crossings times [3] and for sub-Nyquist sampling of single transitions in a square wave [4] have previously been studied. The method

studied in this paper is able to handle multiple square pulses in a single sampling period. We characterize the performance of this method subject to quantization error and integrator noise, in order to understand feasibility of hardware implementation.

Generalized Analog Thresholding (gAT) is one emerging class of method for sampling and reconstructing square waves, derived from the literature on finite rate of innovation, with application to efficient, real-time neurotechnology for basic neuroscience and conditions like epilepsy and paralysis. Methods from the gAT class are denoted gAT- n , where integer $n \geq 1$ is the maximum number of square wave pulses within the analysis interval. The sampling process is described in Section 2.2, and the reconstruction process is described in Section 2.3.

In one possible application of gAT to neurotechnology, a dense, implanted intracortical electrode array records action potentials (spikes) from thousands of neurons simultaneously, monitoring for pre-seizure activity. Each channel multiplexes through a comparator to determine if the voltage from the channel passes a selected threshold [5]. The output of the comparator appears as a sequence of square waves, which is the input signal for gAT. The goal is for gAT to reliably sample and reconstruct these square wave signals at tens of hertz, where spiking neural activity is routinely sampled at 5 – 30 kHz.

The goal of efficient, large-scale spike sampling and reconstruction is made difficult by noise sources and quantization error. Previous work has attempted to deal generally with noise in FRI methods [6, 7, 8]. In this paper, we focus on hardware-related non-idealities that are specific to gAT, including integrator noise and quantization error. We explore three orders of integrator noise and realistic square wave amplitudes based on the IVC102 integrator chip [9]. We chose the IVC102 primarily as a first-pass example with readily available specifications rather than as a canonical low-power chip. Our model for the input signal, sampling, and reconstruction is described in Section 2. The effects of integrator noise and quantization error are systematically explored in Section 3.

2. SIMULATION MODEL

This section describes our model to evaluate hardware feasibility of gAT. The input signal model is described in Section 2.1. The sampling stage, including integrator noise and ADC quantization error, is explained in Section 2.2. The reconstruction process is described in Section 2.3. Model parameter values for simulation are documented in Section 2.5.

2.1. Input Signal

The signal for each sampling interval consists of a sequence of square waves. It is possible for a sequence to consist of zero square waves. A possible signal is shown in Figure 1.




Fig. 1. An example signal for the sampling interval from $t = 0$ to $t = T$. This signals contains two square waves. The first pulse has a duration of w_1 and is centered at time t_1 , and the second pulse has a duration of w_2 and is centered at time t_2 .

More formally, suppose that a sampling interval of length T consists of n square waves, centered on t_1, \dots, t_n with widths w_1, \dots, w_n . The values are valid only when $0 \leq t_1 - \frac{w_1}{2}$, $t_i + \frac{w_i}{2} < t_{i+1} - \frac{w_{i+1}}{2}$ for $i \in \{1, \dots, n-1\}$, and $t_n + \frac{w_n}{2} \leq T$. In other words, the square waves are in order and do not overlap. The signal is then defined as:

$$x(t) = \begin{cases} 1 & t_1 - \frac{w_1}{2} \leq t \leq t_1 + \frac{w_1}{2} \\ \vdots & \vdots \\ 1 & t_n - \frac{w_n}{2} \leq t \leq t_n + \frac{w_n}{2} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

2.2. gAT Sampling Stage with Quantization Error and Integrator Noise

The analog preprocessing stage for gAT- n computes the first $2n$ repeated integrals of the signals, evaluated from the start to the end of the sampling interval. The continuous-time behavior of this stage is evaluated with exact integration, rather than by numerical integration in discrete time.

Under noise-free conditions, the gAT samples are denoted y_1, \dots, y_{2n} , where

$$x_1(t) = \int_0^t x(\tau) d\tau \quad (2)$$

$$x_{k+1}(t) = \int_0^t x_k(\tau) d\tau \quad \text{for } k \in \{1, \dots, 2n-1\} \quad (3)$$

$$y_k = x_k(T) \quad \text{for } k \in \{1, \dots, 2n\} \quad (4)$$

The y_k are then quantized through an ADC to allow reconstruction of $t_1, \dots, t_n, w_1, \dots, w_n$ on a digital system. The n -th order integration motif was previously described for sampling in FRI [10].

The full gAT sampling stage including integrator noise and quantization error is shown in Figure 2. First, the α scales the continuous-time square-wave input $x(t)$. Next, a serial bank of noisy integrators introduces $w_1(t), \dots, w_{2n}(t)$, which are independent, identically-distributed zero-mean continuous-time white Gaussian noise processes. Integrator outputs are scaled by β_i to maximally utilize the dynamic range of the ADCs. The

ADC-sampled values y_1, \dots, y_{2n} are passed to the gAT reconstruction stage (Section 2.3). The scaling factors α and β_i are chosen by a process described in Section 2.5.

After integration, gAT samples are quantized through an ADC. In practice, ADCs introduce several non-idealities [11, 12, 13]. Aperture jitter samples the signal at irregular intervals. Non-linear distortion corrupts the analog signal before quantization. Quantization introduces numerical error to varying extents based on bit resolution. Our model for the ADC includes only quantization error.

A corresponding noise model for analog integration consists of continuous-time white Gaussian noise $w_1(t), \dots, w_{2n}(t)$ added to the input of each stage in a sequence of ideal integrators [10]. The integral of $w_1(t)$ from $t = 0$ to $t = 1$ is also Gaussian:

$$\int_0^1 w_1(t) dt \sim \mathcal{N}(0, \sigma^2) \quad (5)$$

The IVC102 data sheet specifies σ for a wide range of input capacitance in the top right panel on page 4, titled ‘‘Total Output Noise vs C_{IN} ’’ [9].




Fig. 2. gAT sampling stage, modeled with integrator noise and quantization error. A square wave input signal $x(t)$ enters a serial bank of noisy integrators and ADCs to obtain the gAT samples y_1, \dots, y_{2n} . Scaling factors α and β_i are chosen to maximize performance under worst-case noise-free conditions (see Section 2.2 for details).

To simulate the system shown in Figure 2, the exact FRI samples y_1, \dots, y_{2n} are computed. Next, the noise caused by a single term w_k at time step k on $(2n - k + 1)$ present and future FRI samples y_k, \dots, y_{2n} is simulated as a $(2n - k + 1)$ -dimensional zero-mean jointly Gaussian random variable z_k , with the following terms in the i -th row and j -th column of its covariance matrix [10]

$$\Sigma_{i,j} = \frac{\sigma^2}{i+j-1} T^{(i+j-1)} \quad (6)$$

In the noisy integrator model, the i -th noisy FRI sample \tilde{y}_i is simulated as the sum of the i -th noise-free FRI sample (equation 4), and the corresponding $(k - i + 1)$ -th dimension of each multivariate random variable z_1, \dots, z_k . See [10], p. 48 for an analytical expression of the covariance matrix corresponding to this procedure.

2.3. gAT Reconstruction Stage for Square-Wave Pulses

The reconstruction process takes FRI samples y_1, \dots, y_{2n} and computes estimates of the pulse times and pulse widths. For

gAT-1, a closed-form solution can be computed, disregarding non-idealities. Assuming only one square wave pulse occurs in the interval, centered at time t_1 and width w_1 ,

$$x_1(t) \equiv \int_0^t x(\tau) d\tau \quad (7)$$

$$= \begin{cases} 0 & t < t_1 - \frac{w_1}{2} \\ t - (t_1 - \frac{w_1}{2}) & t_1 - \frac{w_1}{2} \leq t \leq t_1 + \frac{w_1}{2} \\ w_1 & t_1 + \frac{w_1}{2} < t \end{cases} \quad (8)$$

$$x_2(t) \equiv \int_0^t x_1(\tau) d\tau \quad (9)$$

$$= \begin{cases} 0 & t < t_1 - \frac{w_1}{2} \\ \frac{1}{2}[t - t_1 + \frac{w_1}{2}]^2 & t_1 - \frac{w_1}{2} \leq t \leq t_1 + \frac{w_1}{2} \\ w_1(t - t_1) & t > t_1 + \frac{w_1}{2} \end{cases} \quad (10)$$

$$y_1 \equiv x_1(T) = w_1 \quad (11)$$

$$y_2 \equiv x_2(T) = w_1(T - t_1) \quad (12)$$

Reordering these terms, the parameters t_1 and w_1 of the pulse are estimated as:

$$\hat{w}_1 = y_1 \quad (13)$$

$$\hat{t}_1 = T - \frac{y_2}{y_1} \quad (14)$$

For gAT-2, we also determined a closed-form solution, not included here due to space constraints. For $n > 2$, the solution may be computed numerically, not explicitly confirmed or explored in this work. The reconstructed pulse parameters were computed using standard MATLAB double precision variables.

2.4. Metrics to Evaluate Signal Reconstruction Quality

We examined gAT performance on the basis of error in the reconstructed pulse time and width, defined as follows. Denote t as the true pulse time, \hat{t} as the reconstructed pulse time, w as the true width, and \hat{w} as the reconstructed width. The mean unsigned error for the reconstructed time is $\mathbb{E}(|\hat{t} - t|)$. The mean unsigned error for the reconstructed width is $\mathbb{E}(|\hat{w} - w|)$. Both are non-negative as the absolute value of the signed error.

2.5. Ranges of Model Parameter Values Explored in Simulation

Input square wave signals are generated as follows. Each sampling interval for gAT-1 contains one randomly generated pulse. Each sampling interval for gAT-2 contains two randomly generated pulses. The time of each pulse is selected uniformly from the entire sampling interval, while ensuring pulses are non-overlapping.

Pulse width w is drawn from a log-normal distribution $\ln \mathcal{N}(\mu = -9, \sigma = 1.15)$. This generates typical pulse widths between 0.013 ms and 1.12 ms. These values are selected to model the distribution of the time a neural spike exceeds a comparator threshold in the analog thresholding approach [5]. Action potentials usually in range of 1 ms width, but the time spent above comparator threshold is considerably shorter.

The integrator noise parameter σ from Equation 5 and scaling parameters α and β_i from Figure 2 are selected based on an example commercial integrator, the Texas Instruments Precision Switched Integrator Transimpedance Amplifier (IVC102)

[9]. The data sheet specifies output noise, input current range and output voltage range. The root mean squared output noise ranges from $4 \mu\text{V}$ to $300 \mu\text{V}$ for a 1 ms integration time. Under the assumption that the integrator output noise scales linearly with integration time, the standard deviation of the output ranges from $4\sqrt{T} \mu\text{V}$ to $300\sqrt{T} \mu\text{V}$. For an interval of 1 second, this standard deviation is the σ and ranges from 0.1 mV to 10 mV.

With regards to input and output ranges, the IVC102 is a transimpedance amplifier, which means that it integrates input current and outputs voltage, as specified in the data sheet:

$$V_{\text{OUT}} = \frac{-1}{C_{\text{INT}}} \int I_{\text{IN}}(t) dt \quad (15)$$

The input range is $\pm 100 \mu\text{A}$ and the output range is $\pm 10 \text{V}$. Based on equation 15, an input of $10 \mu\text{A}$ would saturate the output within 100 microseconds. Because typical pulse widths for the neuroscience application are on the order of 0.013 to 1.12 ms, the input range is more than sufficient to saturate the output. As such, the output range is the limiting variable in choosing α and β_i .

Specifically, the scaling factor α is chosen to avoid IVC102 output saturation while maximally utilizing the dynamic range of the integrator output relative to integrator noise. The scaling factors β_i are chosen to scale the integrator output to match the ADC input range. These values are computed separately for gAT-1 and gAT-2.

First, the α is chosen based on square wave inputs located at the beginning of the analysis interval, resulting in the largest integrals. These square wave inputs vary from 0 to 1, as depicted in Figure 1. This selection disregards the effects of integrator noise. Next, the β_i are chosen so that the ADC input range covers the 95 percent confidence interval of possible samples, accounting for integrator noise. We have not included the α and β_i values because there are several chosen values per figure depending on interval length (affecting α and β_i) and noise level (affecting β_i). The ADC quantization levels are evenly spaced throughout this interval.

Table 1. Integration Noise and Sampling Parameters Explored in Figures 3, 4 and 5.

Figure	σ (mV)	ADC Res. (Bits)	Samp. Period (ms)
3	0 – 10	16	100
4	0, 0.1, 10	1 – 16	100
5	0, 0.1, 10	16	0 – 100

Table 1 lists the integrator noise standard deviation (σ) in mV (reported in the IVC102 data sheet [9] as standard deviation of the output after integrating zero input for 1 second), the ADC resolution in bits, and the length of the sampling period in ms.

3. RESULTS

First, we assessed performance versus integrator noise for a high resolution (16-bit) ADC (Figure 3) and 100 ms sampling interval. Errors grow with integrator noise for both methods as the noise level increases, but much more rapidly for gAT-2 versus gAT-1. This is likely because gAT-2 depends on higher order integrals that accumulate larger variability than lower order

integrals in the presence of integrator noise. The gAT-2 performance degradation asymptotes because pulse time and width reconstruction errors are limited by the values represented by the ADC and the length of the sampling interval.




Fig. 3. Using a high-resolution ADC (16 bit), signal reconstruction quality as a function of the integrator noise level σ . Sampling interval is 100 ms. Typical simulated spike widths vary between 0.013 ms and 1.12 ms. Confidence intervals (95%) are based on 100 bootstrapped averages from 10,000 samples.

We then determined performance versus ADC bit resolution at three levels of integrator noise (Figure 4) for a 100 ms sampling interval, based on the IVC102 integrator [9]. Solid curves for each method represent performance at two different levels of integrator noise (σ in equation 5), 0.1 mV and 10 mV. Dotted curves represent performance at zero integrator noise. Note that all three gAT-1 (blue) curves overlap in this figure. As the number of bits used by the ADC increases, the error decreases because the values received from the ADC become more accurate. In general, gAT-2 has larger errors than gAT-1, suggesting its reconstruction process is more sensitive to errors in the samples. The greatest benefit in ADC resolution is achieved within 8 bits. For gAT-2 curves, unusual behavior is seen at resolution below 3 bits, which likely relates to the relationship between reconstruction errors and the specific conditions of our performance testing.




Fig. 4. Signal reconstruction quality as a function of the ADC resolution in bits, in the presence of integrator noise and quantization error. The solid lines show the expected error for a low level of integrator noise ($\sigma = 0.1$ mV) and a high level of integrator noise ($\sigma = 10$ mV) based on the IVC102 [9]. The dashed lines show performance with no integrator noise, $\sigma = 0$ mV. For gAT-1, the various lines overlap. Confidence intervals (95%) are based on 100 bootstrapped averages from 10,000 samples.

Finally, we examined performance versus sampling interval (Figure 5) at varying levels of integrator noise using a 16-bit ADC. As the sampling interval increases in duration, the error in reconstructed time increases, due to accumulated integrator

noise. Paradoxically, the reconstructed width error decreases with increasing sampling interval. The reason for this is not immediately clear.

For practical applications, an additional source of performance degradation relates to the probability of exceeding the assumed maximum number of pulses within an interval. For example, neuron spiking rates vary between 1 and 300 action potentials per second depending on species, brain region, sensory stimulation, and other conditions [14]. Accordingly, specific choices of sampling interval might work well in some brain regions and not others. This source of error is not represented in Figure (5).




Fig. 5. Signal reconstruction quality versus length of the sampling interval, using a 16-bit ADC at various levels of integration noise. The solid lines show the expected error for low ($\sigma = 0.1$ mV) and high ($\sigma = 10$ mV) integrator noise. The dashed lines show performance without integration noise. For gAT-1, the lines overlap. Confidence intervals (95%) are based on 100 bootstrapped averages from 10,000 samples.

4. CONCLUSION

Our simulated analysis examines the feasibility of gAT, a class of FRI-based sampling and reconstruction methods for square-wave signals, under the hardware-induced non-idealities of integrator noise and quantization error. Under ideal conditions, FRI-based gAT methods are capable of reconstructing signals precisely. By simulating these hardware non-idealities, we show that gAT-1 reconstruction could be robust to these hardware non-idealities, where gAT-2 reconstruction is expected to be significantly more brittle in terms of reliably estimating pulse time and width. The brittle performance of gAT-2 is likely in part because gAT-2 reconstruction method requires higher order integrals that accumulate noise from lower-order integration stages. Based on these results, gAT-1 appears to be the preferred candidate over gAT-2 for initial efforts to develop FRI sampling hardware for square-wave signals.

5. ACKNOWLEDGMENTS

The authors would like to thank Julius Kusuma (Schlumberger-Doll Research), Lav Varshney (Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign), and Juhwan Yoo (Broadcom) for discussions on FRI and hardware implementation. B.D.H. and A.W. were supported by the Caltech Summer Undergraduate Research Fellowship. L.S. was supported by the American Heart Association Scientist Development Grant (11SDG7550015).

6. REFERENCES

- [1] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, 2006.
- [2] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *Signal Processing, IEEE Transactions on*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [3] C. S. Seelamantula, "A sub-nyquist sampling method for computing the level-crossing-times of an analog signal: Theory and applications," in *Signal Processing and Communications (SPCOM), 2010 International Conference on*, pp. 1–5, IEEE, 2010.
- [4] G. Ramesh, E. Atallah, and Q. Sun, "Recovery of bilevel causal signals with finite rate of innovation using positive sampling kernels," pp. 129–132, 2013.
- [5] R. R. Harrison, "The design of integrated circuits to observe brain activity," *Proceedings of the IEEE*, vol. 96, no. 7, pp. 1203–1216, 2008.
- [6] A. Wein and L. Srinivasan, "Iterml: A fast, robust algorithm for estimating signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 61, pp. 5324–5336, 2013.
- [7] T. Blu, P.-L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 31–40, 2008.
- [8] V. Y. F. Tan and V. K. Goyal, "Estimating signals with finite rate of innovation from noisy samples: A stochastic algorithm," *Signal Processing, IEEE Transactions on*, vol. 56, no. 10, pp. 5135–5146, 2008.
- [9] Texas Instruments, *Precision Switched Integrator Transimpedance Amplifier*, 2000. Part #: IVC102.
- [10] J. Kusuma, *Economical sampling of parametric signals*. PhD thesis, Massachusetts Institute of Technology, 2006.
- [11] C. R. Parkey, W. B. Mikhael, D. B. Chester, and M. T. Hunter, "Modeling of jitter and its effects on time interleaved adc conversion," in *AUTOTESTCON, 2011 IEEE*, pp. 367–372, IEEE, 2011.
- [12] C. R. Parkey, D. B. Chester, M. Hunter, and W. Mikhael, "Simulink modeling of analog to digital converters for post conversion correction development and evaluation," in *Circuits and Systems (MWSCAS), 2011 IEEE 54th International Midwest Symposium on*, pp. 1–4, IEEE, 2011.
- [13] L. Michaeli and J. Šaliga, "Error models of the analog to digital converters," *Measurement Science Review*, vol. 14, no. 2, pp. 62–77, 2014.
- [14] E. R. Kandel, J. H. Schwartz, T. M. Jessell, *et al.*, *Principles of neural science*, vol. 4. McGraw-Hill New York, 2000.