Hardware-Limited Time Constant Estimation Using a Weighted Linear Regression

Titan Yuan, Filip Maksimovic, David C Burnett, Kristofer S.J. Pister

HARDWARE-LIMITED TIME CONSTANT ESTIMATION USING A WEIGHTED LINEAR REGRESSION

Titan Yuan[⋆]*, Filip Maksimovic*[†]*, David C. Burnett*[‡]*, Kristofer S.J. Pister*[⋆]

[⋆]University of California, Berkeley, CA [†]Inria, Paris [‡]Portland State University, Portland, OR

ABSTRACT

Accurately determining the time constant of a circuit enables IoT nodes to easily read out resistive or capacitive sensors. However, power and cost constraints lead to hardware limitations that complicate such measurements, including ADC noise, sampling clock jitter, poor voltage control over temperature and process, and a low-power microprocessor without a fast multiplier or floating point support. This work discusses estimating the time constant of a decaying exponential's ADC samples using a simple weighted linear regression and describes the on-chip implementation of the regression on a low-cost, low-power microprocessor. Experimental results with an imperfect ADC show that time constants over more than two orders of magnitude can be accurately estimated within 5% of the nominal value with a mean standard error of about 1% of the nominal value.

***Index Terms*—** Weighted linear regression, time constant estimation, resistive sensor readout, internet of things, low-cost sensing

1. INTRODUCTION

With the growing popularity of the internet of things and wearable electronics [1], low-cost and low-power sensing with compact readout electronics is becoming increasingly essential. In particular, resistive sensors are cost-effective at measuring temperature [2, 3], gas concentrations [4, 5], strain [3], and magnetic fields [6]. However, low-power ADCs on low-cost microcontrollers suffer from noisy ADC samples, variable LSB step sizes, and poorly controlled reference voltages, if any. As a result, a resistive divider is infeasible without prior multi-point voltage calibration.

Instead, a reference-independent method to measure the resistance is to connect a known capacitor with an initial voltage in parallel to the resistor under test, as shown in Fig. 1(a), and estimate the circuit's time constant by sampling the exponentially decaying voltage and post-processing the samples. While the accuracy of the estimated time constant is independent of the reference capacitor, the accuracy of the estimated resistance is directly proportional to that of the reference capacitor. However, capacitors with a tolerance of $<2\%$ and high temperature stability (class I capacitors) can be easily acquired and are much cheaper than an off-chip reference voltage.

We would like to emphasize that although resistive sensing is the primary motivation for this work, our main contribution is in estimating the time constant of a general decaying exponential. Section 3 proposes processing the exponentially decaying ADC samples with a weighted linear regression, for which the weights are derived in the appendix. In Section 4, we describe the shortcomings of the microcontroller's ADC hardware in more detail, and Section 5 presents some experimental results on estimating the time constant using the microcontroller's imperfect ADC.

This research was funded by ARPA-E under award number DE-AR0001602. Corresponding author: Titan Yuan (titan@berkeley.edu).

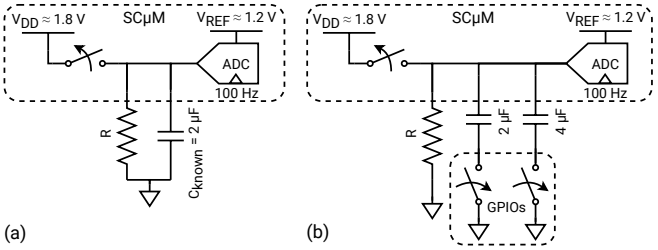

Fig. 1. (a) Test circuit schematic to determine the unknown resistance R. C_{known} is initially charged to V_{DD} by a GPIO pin before the switch is opened and the ADC starts sampling the voltage across R and C_{known} at $f_s = 100\,\mathrm{Hz}$. (b) The dynamic range of the measurable resistance can be extended with a simple capacitive DAC.

2. PRIOR WORK

Typically, front-end circuits for resistive sensing rely on a resistance-to-voltage or a resistance-to-time conversion. The Wheatstone bridge [7], a voltage divider, and a transresistance amplifier [3] are examples of the former, but they require a known constant supply voltage, or current source, and an accurate ADC to measure the output voltage, all of which are not readily available on a low-cost microcontroller. Furthermore, the resistance resolution varies due to the nonlinear output voltage, and the measurable resistance range must be known in advance.

Resistance-to-time conversion methods are usually able to measure a wide range of resistances by using a relaxation oscillator, where the resistor current repeatedly charges and discharges a capacitor [8]. [9] proposes a compensation function based on the high and low periods to derive a voltage-independent resistance estimate.

Finally, [10] discharges a known capacitor across the unknown resistance and measures the time at which the exponential decays past a certain threshold. Any parasitic lead resistance is compensated for by discharging over different combinations of resistances. However, estimating the resistance with only the time of decay is insufficiently accurate, especially with a noisy ADC, and our resistances under test are much larger than the parasitics.

[11] uses the same chip as this paper to measure the mean H_2S sensor readout with the chip's ADC over 200 samples and transmit it wirelessly at $2.4\,\mathrm{GHz}$. Averaging reduces the sample variance but requires a constant input during the averaging period, which does not hold for an exponentially decaying voltage.

3. TIME CONSTANT ANALYSIS

Given the ADC samples of a decaying exponential, we propose a weighted linear regression approach to estimate its time constant.

We model the kth ADC sample of the decaying exponential as follows:

$$s[k] = ae^{-\frac{k}{f_s \tau}} + n_k, n_k \sim \mathcal{N}(0, \sigma^2), \tag{1}$$

where a is the exponential's scaling factor, f_s is the sampling rate, τ is the unknown time constant, and n_k models the ADC noise. In reality, $n_k = w_k + q_k$ consists of some Gaussian noise $w_k \sim \mathcal{N}(0, \sigma_w^2)$ and some uniformly distributed quantization noise q_k with $\sigma_q = \frac{1}{\sqrt{12}}$ LSB. Since $\sigma_w \gg \sigma_q$ for a noisy ADC, we model n_k as a single Gaussian, which turns out to be a valid assumption.

As derived in (6), we take the natural logarithm to obtain a linear model:

$$\ln(s[k]) \approx \ln(a) - \frac{1}{\tau}\frac{k}{f_s} + \tilde{n}_k, \tag{2}$$

where $\tilde{n}_k \sim \mathcal{N}(0, \tilde{\sigma}_k^2)$ and $\tilde{\sigma}_k^2 = \sigma^2/a^2 \cdot \exp\left(\frac{2k}{f_s \tau}\right)$ according to (17). The approximation is justified by Fig. 2, where we plot the theoretical standard deviation based on (13), the approximated standard deviation based on (17), and the standard deviation based on simulated ADC data.

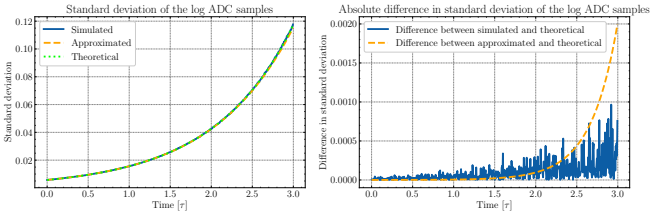

Fig. 2. Theoretical, approximated, and simulated distribution of the log ADC data samples.

Since the sample variance in log space is no longer equal for all samples, we propose a weighted linear regression as follows:

1. Determine a rough estimate of 3τ by finding the first ADC sample at which the exponential has decayed by 95%. Estimate a as the first ADC sample, i.e., when $k = 0$.
2. Using only the ADC samples through the 3τ index, perform a weighted linear regression on $\ln(s[k])$ with the weight of the kth sample given by $w_k = a^2/\sigma^2 \cdot \exp\left(-\frac{2k}{f_s \tau}\right)$, as derived in (18).

The weights are exponentially decreasing because the logarithm maps the variance at larger ADC values, or smaller ADC sample indices k, to a smaller variance, as shown in Fig. 3. Similarly, only ADC samples through 3τ are used because for larger k, $\ln(s[k])$ is dominated by noise, whose quantization noise component is additionally no longer uniformly distributed. It is important to also emphasize that step 1 only provides a rough estimate of τ, usually underestimating it due to the presence of noise, but the estimate of τ is a sufficiently good heuristic for determining the weights.

In step 2, the estimated slope $-1/\hat{\tau}$ has a closed-form solution that can be derived from the normal equations [12]:

$$-\frac{1}{\hat{\tau}} = \frac{\sum_k w_k \left(\frac{k}{f_s} - \bar{x}\right)(\ln(s[k]) - \bar{y})}{\sum_k w_k \left(\frac{k}{f_s} - \bar{x}\right)^2}, \tag{3}$$

where

$$\bar{x} = \frac{\sum_k w_k \frac{k}{f_s}}{\sum_k w_k}, \bar{y} = \frac{\sum_k w_k \ln(s[k])}{\sum_k w_k}. \tag{4}$$

The variances of the estimated slope $1/\hat{\tau}$ and of the corresponding time constant $\hat{\tau}$ are approximated by (30) and (31), respectively.

To evaluate the proposed weighted linear regression, we compared it with an exponential regression, a normal linear regression, and a polynomial regression. The exponential regression was performed using `scipy`'s optimizer as it is a nonlinear optimization problem. We simulated a decaying exponential with a vertical offset of 127 LSBs, an ADC noise of $\sigma = 5$ LSBs, a sampling rate $f_s = 100$ Hz, and a time constant of $\tau = 2$ s. The result is shown in Fig. 3. Clearly, the rough estimate of τ based on 95% decay underestimates τ.

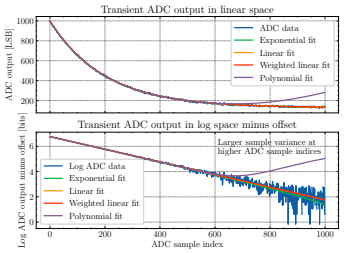

Regression type	Estimated τ [s]
Nominal	2
Based on 95% decay	1.897
Exponential fit	2.003
Linear fit	1.993
Weighted linear fit	1.998
Polynomial fit (degree 5)	2.013

Fig. 3. Fitting a simulated decaying exponential ($\tau = 2$ s) sampled at $f_s = 100$ Hz with different regression types.

4. ADC HARDWARE LIMITATIONS

For this paper, we estimate the time constant using the on-chip 10-bit sensor ADC of the single-chip micro mote (SCµM), a $2 \times 3 \times 0.3$ mm^3 low-cost, low-power, and crystal-free system-on-chip an ARM Cortex M0 microprocessor, 64 kB of data and instruction memory each, and a 2.4 GHz transceiver for IoT and microrobotic applications [13]. Its 10-bit sensor SAR ADC, based on [14], is intended to interface with external signals but has multiple shortcomings due to low SWaP-C constraints.

The ADC has about $\sigma = 5$ LSBs of noise, and the ADC samples tend to be "sticky" towards larger powers of 2 due to non-monotonicities in the binary DAC causing peaks in the DNL of the ADC [15]. Fig. 4 shows a histogram of the ADC samples at a fixed voltage resulting from this "stickiness." Furthermore, the ADC sampling timer is derived from SCµM's 20 MHz CPU clock, which has a measured RMS noise of at least 500 ppm [16] and thus already reduces the ADC's expected SNR by -30 dB at $f_s = 100$ Hz [17].

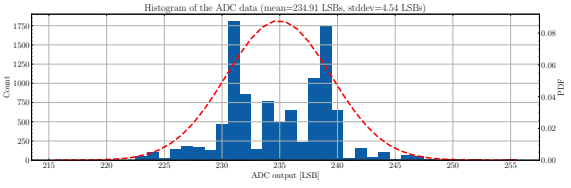

Fig. 4. Histogram of the ADC samples at a fixed voltage overlaid with the normal distribution expected from a typical ADC. The modes correspond to 231 ($= 255 - 16 - 8$) and 239 ($= 255 - 16$) LSBs.

The MSB of the 10-bit ADC is also always low due to a bug in the ADC's RTL, so the ADC samples must be corrected prior to processing, as shown in Fig. 5. However, since $V_{\text{DD}} > V_{\text{REF}}$, we know *a priori* that the MSB starts equal to 1, i.e., the ADC output should be between 512 and 1023 LSBs.

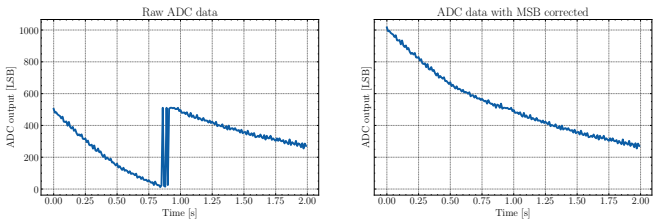

Fig. 5. ADC data before and after correcting the MSB.

Lastly, SCμM has poor voltage control, and the on-chip regulated ADC reference varies between 1.2 V and 1.3 V among chips. SCμM's ADC also has a non-zero offset, such that a 0 V input does not result in an ADC output of 0 LSBs. Therefore, to estimate the zero offset, we first sample the decaying exponential until the running minimum ADC output has not changed within the last 1.5 s. We then estimate the zero offset as the average of the last hundred ADC samples and subtract this offset from all ADC samples.

With these ADC idiosyncrasies in mind, we used SCμM to collect actual ADC data of multiple decaying exponentials with varying time constants. As shown in Fig. 1(a), we used a fixed 2 μF 50 V ceramic capacitor and varied the resistance from 10 kΩ up to 10 MΩ with 1% tolerance 1 W metal film resistors. The value of the capacitor was verified using a Fluke 289 multimeter while the values of the resistors were verified using a four-probe measurement on a Keithley 2400-LV sourcemeter. We then compared the performance of the different regression types, as shown in Fig. 6.

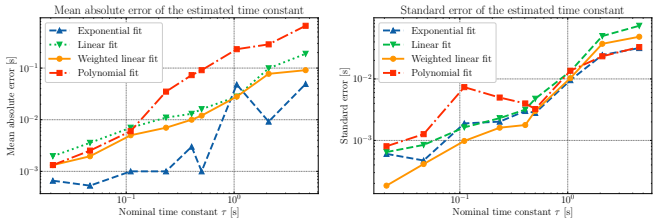

Fig. 6. Comparing the mean absolute error and standard error of the estimated time constant for different regression types. The ADC data was collected according to Fig. 1(a) and processed offline.

While an exponential regression has a lower estimation variance than a linear regression, the nonlinear optimization is difficult to implement on a microprocessor. However, a weighted linear regression performs better than an ordinary one with a comparable mean absolute error but lower estimation variance; the R^2 value for a linear regression in log space is around 0.998 while the R^2 value for a weighted linear regression is greater than 0.9999.

5. EXPERIMENTAL RESULTS

We decided to process the ADC samples on-chip, characteristic of edge computing. This primarily avoids occupying valuable network bandwidth with raw ADC data and minimizes power consumption due to a shorter radio TX on time. Thus, we implemented the algorithm described in Section 3, including the reciprocal of (3), in SCμM's firmware using a sampling frequency of $f_s = 100\,\text{Hz}$.

However, SCμM's ARM Cortex M0 presents some limitations: First, due to a lack of a floating-point unit, float operations are expensive, and even simple operations, e.g., converting an integer to a float, can cause SCμM's CPU to suddenly crash. As a result, to avoid all floats, we used Q50.14 fixed-point integers. Since the ADC outputs 10-bit integers, we used a lookup table with 1024 values to calculate $\ln(\cdot)$, and $e^{(\cdot)}$ was calculated using a second-order Taylor series approximation. Second, multiplying 64-bit integers would lead to CPU crashes as well, so multiplication was done by adding bit-shifted copies of one of the factors. The root cause of these crashes has yet to be determined. Finally, since SCμM has a limited 64 kB data memory, we limited the ADC sample buffer to 5000 samples.

Fig. 7 shows the results of running the on-chip implementation on SCμM. The output of SCμM is a fixed-point integer that we divided by 2^{14} offline to obtain the estimated time constant. Note that the simulated and the approximated standard error match well, justifying the approximations made in Section 7.2. The difference between the measured and the simulated standard error is due to unmodeled sources of variance, primarily the zero offset estimate of the decaying exponential and possibly some clock jitter.

Nevertheless, between $\tau = 0.02\,\text{s}$ and $\tau = 10\,\text{s}$, more than two orders of magnitude, the mean estimated time constant was within 5% of the nominal time constant with a standard error of around 1% of the nominal time constant. The minimum measurable time constant is limited by the MSB disambiguation algorithm because the exponential decays too fast to correctly determine the MSB. In our experiment with $f_s = 100\,\text{Hz}$, the minimum estimatable time constant was around $\tau = 0.02\,\text{s}$, which could still be accurately estimated with a 3τ index of 7 ADC samples. The maximum measurable time constant is constrained by the on-chip memory buffer size given the sampling frequency.

Runtime-wise, SCμM usually gathers around 5τ's worth of ADC samples, of which only around 60% are processed. The sampling period could be shortened by a dynamic threshold similar to [18]. On SCμM, processing around 290 samples took about 300 ms while processing around 3600 samples took about 20 s. Note that the fixed capacitor is charged for every measurement, so the energy per measurement is proportional to the capacitance.

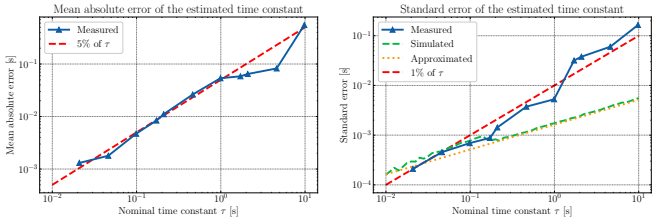

Fig. 7. Mean absolute error and standard error of the estimated time constant when collecting the ADC data according to Fig. 1(a) and processing it on-chip with the fixed-point implementation of the weighted linear regression. The approximated standard error is given by (31).

6. CONCLUSION

To estimate a circuit's time constant while sampling its exponential decay, we derived an approximation for the sample variance in log space and proposed using a weighted linear regression on the first 3τ's worth of ADC samples. This time constant analysis method is independent of the ADC's reference voltage and LSB step size and can be implemented easily and with reasonable runtime even on a processor with no floating point support. Experimental results on SCμM with a single fixed capacitor showed that time constants over

more than two orders of magnitude could be estimated within 5% of the nominal value with a mean standard error of about 1% of the nominal value. The actual dynamic range of measurable resistances depends on the ADC sampling frequency but can be extended with a simple capacitive DAC, as shown in Fig. 1(b).

Although experiments were only performed on SCµM, the proposed weighted linear regression method discussed in Section 3 can be applied to any low-power ADC. Enabling hardware-constrained system-on-chips to easily and accurately estimate the time constant over multiple orders of magnitude, even with a subpar ADC, will allow them to determine unknown resistances or capacitances and operate as small, low-cost, and network-enabled sensors for IoT applications.

7. APPENDIX

7.1. Derivation Of The Weights

Taking the natural logarithm of (1), we obtain a linear model in k:

$$\ln(s[k]) = \ln(ae^{-\frac{k}{f_s\tau}} + n_k) \tag{5}$$

$$= \ln(a) - \frac{k}{f_s\tau} + \ln\left(1 + \frac{n_k}{a}e^{\frac{k}{f_s\tau}}\right) \tag{6}$$

Let $X_k = \ln\left(1 + \frac{n_k}{a}e^{\frac{k}{f_s\tau}}\right)$ denote the noise term of the kth sample after the logarithm. Finding its CDF and PDF, we get:

$$F_{X_k}(x) = \Pr\{X_k \le x\} \tag{7}$$

$$= \Pr\left\{n_k \le ae^{-\frac{k}{f_s\tau}}(e^x - 1)\right\} \tag{8}$$

$$= F_{n_k}\left(ae^{-\frac{k}{f_s\tau}}(e^x - 1)\right) \tag{9}$$

$$f_{X_k}(x) = F'_{X_k}(x) \tag{10}$$

$$= f_{n_k}\left(ae^{-\frac{k}{f_s\tau}}(e^x - 1)\right) \cdot ae^{x-\frac{k}{f_s\tau}} \tag{11}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{a^2e^{-\frac{2k}{f_s\tau}}(e^x-1)^2}{2\sigma^2}} \cdot ae^{x-\frac{k}{f_s\tau}} \tag{12}$$

Define a new variable $l = \frac{k}{f_s\tau}$.

$$f_{X_k}(x) = \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{a^2e^{-2l}(e^x-1)^2}{2\sigma^2}} \cdot ae^{x-l} \tag{13}$$

Plotting f_{X_k} at various values of τ shows that the support of X_k is fairly narrow about $x = 0$. This observation allows us to approximate $e^x \approx 1 + x$, which could be further approximated as $1 + x \approx 1$.

$$f_{X_k}(x) \approx \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{a^2e^{-2l}x^2}{2\sigma^2}} \cdot ae^{-l}(1 + x) \tag{14}$$

$$\approx \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{a^2e^{-2l}x^2}{2\sigma^2}} \cdot ae^{-l} \tag{15}$$

$$= \frac{1}{\sqrt{2\pi}\tilde{\sigma}}e^{-\frac{x^2}{2\tilde{\sigma}^2}}, \tag{16}$$

which we recognize as the PDF of a Gaussian random variable with variance:

$$\tilde{\sigma}^2 = \frac{\sigma^2}{a^2}e^{2l} = \frac{\sigma^2}{a^2}e^{\frac{2k}{f_s\tau}}. \tag{17}$$

Therefore, we set the weight of the kth log sample in the weighted linear regression as:

$$w_k = \frac{1}{\tilde{\sigma}^2} = \frac{a^2}{\sigma^2}e^{-2l} = \frac{a^2}{\sigma^2}e^{-\frac{2k}{f_s\tau}}. \tag{18}$$

7.2. Variance Of Estimated Time Constant

From (3), since w_k contains all of the sample variance, $\mathrm{Var}(\ln(s[k]) = 1$, so the variance of the estimated slope is:

$$\mathrm{Var}\left(-\frac{1}{\hat{\tau}}\right) = \mathrm{Var}\left(\frac{1}{\hat{\tau}}\right) = \frac{1}{\sum_k w_k\left(\frac{k}{f_s} - \bar{x}\right)^2} \tag{19}$$

To find $\mathrm{Var}(\hat{\tau})$, we perform multiple approximations, justfied by the fact that if $f_s = 100\,\mathrm{Hz}$, $f_s\tau \gg 2 \implies e^{-\frac{2}{f_s\tau}} \approx 1 - \frac{2}{f_s\tau} \approx 1$.

$$\sum_k w_k = \frac{a^2}{\sigma^2}\sum_k e^{-\frac{2k}{f_s\tau}} \approx \frac{a^2}{\sigma^2}\frac{1}{1 - e^{-\frac{2}{f_s\tau}}} \approx \frac{a^2}{\sigma^2}\frac{f_s\tau}{2} \tag{20}$$

Using the fact that $\sum_{x=0}^{\infty} xr^x = \frac{r}{(1-r)^2}$ for $|r| < 1$:

$$\sum_k w_k\frac{k}{f_s} = \frac{a^2}{\sigma^2 f_s}\sum_k e^{-\frac{2k}{f_s\tau}}k \tag{21}$$

$$\approx \frac{a^2}{\sigma^2 f_s}\frac{e^{-\frac{2}{f_s\tau}}}{\left(1 - e^{-\frac{2}{f_s\tau}}\right)^2} \tag{22}$$

$$\approx \frac{a^2}{\sigma^2 f_s}\frac{f_s^2\tau^2}{4} = \frac{a^2}{\sigma^2}\frac{f_s\tau^2}{4} \tag{23}$$

Thus,

$$\bar{x} = \frac{\sum_k w_k\frac{k}{f_s}}{\sum_k w_k} \approx \frac{\frac{a^2}{\sigma^2}\frac{f_s\tau^2}{4}}{\frac{a^2}{\sigma^2}\frac{f_s\tau}{2}} = \frac{\tau}{2} \tag{24}$$

Intuitively, $\bar{x}$ can be viewed as the expected value of an exponential random variable with rate $\frac{2}{\tau}$.

Finally, using $\sum_{x=0}^{\infty} x^2r^x = \frac{r(r+1)}{(1-r)^3}$ for $|r| < 1$:

$$\sum_k w_k\left(\frac{k}{f_s}\right)^2 = \frac{a^2}{\sigma^2 f_s^2}\sum_k e^{-\frac{2k}{f_s\tau}}k^2 \tag{25}$$

$$\approx \frac{a^2}{\sigma^2 f_s^2}\frac{e^{-\frac{2k}{f_s\tau}}\left(e^{-\frac{2}{f_s\tau}} - 1\right)}{\left(1 - e^{-\frac{2}{f_s\tau}}\right)^3} \tag{26}$$

$$\approx \frac{a^2}{\sigma^2 f_s^2}\frac{2f_s^3\tau^3}{8} = \frac{a^2}{\sigma^2}\frac{f_s\tau^3}{4} \tag{27}$$

Plugging (20), (23), (24), and (27) into (19):

$$\mathrm{Var}\left(\frac{1}{\hat{\tau}}\right) = \frac{1}{\sum_k w_k\left(\frac{k}{f_s} - \bar{x}\right)^2} \tag{28}$$

$$= \frac{1}{\sum_k w_k\left(\left(\frac{k}{f_s}\right)^2 - 2\bar{x}\frac{k}{f_s} + \bar{x}^2\right)} \tag{29}$$

$$\approx \frac{1}{\frac{a^2}{\sigma^2}\frac{f_s\tau^3}{4} - \frac{a^2}{\sigma^2}\frac{f_s\tau^3}{4} + \frac{a^2}{\sigma^2}\frac{f_s\tau^3}{8}} = \frac{8\sigma^2}{a^2 f_s\tau^3} \tag{30}$$

For a random variable X, we can approximate $\mathrm{Var}\left(\frac{1}{X}\right) \approx \frac{1}{\mu^4}\mathrm{Var}(X)$, assuming that the support is narrow around μ, by performing a Taylor series approximation of $\frac{1}{X}$ about $X = \mu$.

$$\mathrm{Var}(\hat{\tau}) \approx \tau^4\mathrm{Var}\left(\frac{1}{\hat{\tau}}\right) = \tau^4\frac{8\sigma^2}{a^2 f_s\tau^3} = \frac{8\sigma^2\tau}{a^2 f_s} \tag{31}$$

8. REFERENCES

[1] Wei Gao, Sam Emaminejad, Hnin Yin Yin Nyein, Samyuktha Challa, Kevin Chen, Austin Peck, Hossain M Fahad, Hiroki Ota, Hiroshi Shiraki, Daisuke Kiriya, Der-Hsien Lien, George A Brooks, Ronald W Davis, and Ali Javey, "Fully integrated wearable sensor arrays for multiplexed in situ perspiration analysis," *Nature*, vol. 529, no. 7587, pp. 509–514, Jan 2016.

[2] Chun-Chih Huang, Zhen-Kai Kao, and Ying-Chih Liao, "Flexible Miniaturized Nickel Oxide Thermistor Arrays via Inkjet Printing Technology," *ACS Applied Materials & Interfaces*, vol. 5, no. 24, pp. 12954–12959, 2013.

[3] Pablo Escobedo, Mitradip Bhattacharjee, Fatemeh Nikbakhtnasrabadi, and Ravinder Dahiya, "Smart Bandage With Wireless Strain and Temperature Sensors and Batteryless NFC Tag," *IEEE Internet of Things Journal*, vol. 8, no. 6, pp. 5093–5100, 2021.

[4] Ralf Moos, Noriya Izu, Frank Rettig, Sebastian Reiss, Woosuck Shin, and Ichiro Matsubara, "Resistive oxygen gas sensors for harsh environments," *Sensors (Basel)*, vol. 11, no. 4, pp. 3439–3465, Mar 2011.

[5] JoAnna Milam-Guerrero, Bingxin Yang, Dung T. To, and Nosang V. Myung, "Nitrous Oxide Is No Laughing Matter: A Historical Review of Nitrous Oxide Gas-Sensing Capabilities Highlighting the Need for Further Exploration," *ACS Sensors*, vol. 7, no. 12, pp. 3598–3610, 2022.

[6] Chao Zheng, Ke Zhu, Susana Cardoso de Freitas, Jen-Yuan Chang, Joseph E. Davies, Peter Eames, Paulo P. Freitas, Olga Kazakova, CheolGi Kim, Chi-Wah Leung, Sy-Hwang Liou, Alexey Ognev, S. N. Piramanayagam, Pavel Ripka, Alexander Samardak, Kwang-Ho Shin, Shi-Yuan Tong, Mean-Jue Tung, Shan X. Wang, Songsheng Xue, Xiaolu Yin, and Philip W. T. Pong, "Magnetoresistive Sensor Development Roadmap (Non-Recording Applications)," *IEEE Transactions on Magnetics*, vol. 55, no. 4, pp. 1–30, 2019.

[7] Mingyuan Ren, Honghai Xu, Changchun Dong, and Zhu Zhang, "Toward a Gas Sensor Interface Circuit—A Review," *IEEE Sensors Journal*, vol. 22, no. 19, pp. 18253–18265, 2022.

[8] A. Flammini, D. Marioli, and A. Taroni, "A low-cost interface to high-value resistive sensors varying over a wide range," *IEEE Transactions on Instrumentation and Measurement*, vol. 53, no. 4, pp. 1052–1056, 2004.

[9] K. Elangovan and Chandrika Sreekantan Anoop, "Simple and Efficient Relaxation-Oscillator-Based Digital Techniques for Resistive Sensors — Design and Performance Evaluation," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 9, pp. 6070–6079, 2020.

[10] Ponnalagu Ramanathan Nagarajan, Boby George, and V. Jagadeesh Kumar, "Improved Single-Element Resistive Sensor-to-Microcontroller Interface," *IEEE Transactions on Instrumentation and Measurement*, vol. 66, no. 10, pp. 2736–2744, 2017.

[11] David C. Burnett, Hossain M. Fahad, Lydia Lee, Filip Maksimovic, Brad Wheeler, Osama Khan, Ali Javey, and Kristofer S. J. Pister, "Two-Chip Wireless H2S Gas Sensor System Requiring Zero Additional Electronic Components," in *2019 20th International Conference on Solid-State Sensors, Actuators and Microsystems & Eurosensors XXXIII (TRANSDUCERS & EUROSENSORS XXXIII)*, 2019, pp. 1222–1225.

[12] G. A. F. (George Arthur Frederick) Seber, *Linear regression analysis*, Wiley Series in Probability and Statistics. Wiley-Interscience, Hoboken, N.J, 2nd edition, 2003.

[13] Filip Maksimovic, Brad Wheeler, David C. Burnett, Osama Khan, Sahar Mesri, Ioana Suciu, Lydia Lee, Alex Moreno, Arvind Sundararajan, Bob Zhou, Rachel Zoll, Andrew Ng, Tengfei Chang, Xavier Villajosana, Thomas Watteyne, Ali Niknejad, and Kristofer S. J. Pister, "A Crystal-Free Single-Chip Micro Mote with Integrated 802.15.4 Compatible Transceiver, sub-mW BLE Compatible Beacon Transmitter, and Cortex M0," in *2019 Symposium on VLSI Circuits*, 2019, pp. C88–C89.

[14] M.D. Scott, B.E. Boser, and K.S.J. Pister, "An ultralow-energy ADC for Smart Dust," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 7, pp. 1123–1129, 2003.

[15] Lydia Lee, *Precise Pulse Discrimination for Space-Based Timing Front Ends*, Ph.D. thesis, EECS Department, University of California, Berkeley, Aug 2023.

[16] David C. Burnett, Brad Wheeler, Lydia Lee, Filip Maksimovic, Arvind Sundararajan, Osama Khan, and Kristofer S. J. Pister, "CMOS oscillators to satisfy 802.15.4 and Bluetooth LE PHY specifications without a crystal reference," in *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, 2019, pp. 218–223.

[17] Derek Redmayne, Eric Trelewicz, and Alison Smith, *Understanding the Effect of Clock Jitter on High Speed ADCs*, Linear Technology, 2006.

[18] Andrea De Marcellis, Giuseppe Ferri, Alessandro Depari, and Alessandra Flammini, "A novel time-controlled interface circuit for resistive sensors," in *SENSORS, 2011 IEEE*, 2011, pp. 1137–1140.