

# PANDA: Performance Acceleration through Nonuniform Data Acquisition

Neil Klingensmith  
University of Wisconsin  
naklingensmi@wisc.edu

Suman Banerjee  
University of Wisconsin  
suman@cs.wisc.edu

## ABSTRACT

We present PANDA, a data acquisition technique for energy-constrained sensor nodes that reduces the energy per operation required to sample and preprocess analog sensor data. PANDA takes advantage of the energy consumption patterns of commodity microcontrollers by sampling input signals in short bursts followed by long periods of inactivity. This approach reduces the overhead of repetitively transitioning the CPU and analog components in and out of low-power sleep states. This nonuniformly-spaced input data is then fed to a nonuniform FFT algorithm that computes the frequency spectrum. We show that the spectrum computed with the nonuniform FFT is very close to the spectrum that would be computed from uniformly sampled data preprocessed with a conventional FFT. The output of the nonuniform FFT can be filtered or postprocessed with conventional frequency domain analysis techniques, and a uniformly resampled output can be constructed with the conventional inverse FFT. We compare the energy consumption patterns of burst-mode sampling to those of conventional uniform sampling in several real sensor nodes. We demonstrate that for reasonably sized input datasets, burst mode sampling and postprocessing consumes more than 17% less energy than conventional uniform sampling, including the additional computations required to compute the nonuniform DFT.

## CCS CONCEPTS

• Computer systems organization → Embedded systems;

## KEYWORDS

ADC; Sampling; IoT

## ACM Reference format:

Neil Klingensmith and Suman Banerjee. 2018. PANDA: Performance Acceleration through Nonuniform Data Acquisition. In *Proceedings of The Nineth International Conference on Future Energy Systems, Karlsruhe, Germany, June 12–15, 2018 (e-Energy '18)*, 11 pages. <https://doi.org/10.1145/3208903.3208904>

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*e-Energy '18, June 12–15, 2018, Karlsruhe, Germany*  
© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.  
ACM ISBN 978-1-4503-5767-8/18/06...\$15.00  
<https://doi.org/10.1145/3208903.3208904>

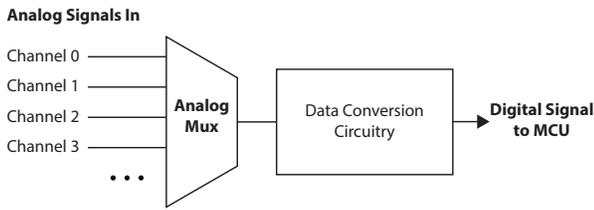
## 1 INTRODUCTION

Energy consumption in mobile and IoT devices is a primary concern because power sources—batteries, energy harvesting devices, etc—are limited in capacity. In ultra-low power sensor systems such as SNUPI [7], MICA2 [8], Amulet [16], and others [2, 18–20], power consumed by analog-to-digital conversion circuitry is a primary contributor to battery drain. As microcontrollers become increasingly complex, adding features such as direct memory access, memory protection, etc., optimizing sleep state power consumption will become more important because wake state power draw will be higher. With PANDA, we aim to reduce the overall power consumption of battery-powered sensing devices by increasing the fraction of time the system spends in a low-power sleep state. PANDA targets applications that use mid-range microcontrollers which are sensitive to power consumption and also have sufficient resources to perform nontrivial computations. This mid-range segment, represented by 32-bit devices such as the ARM Cortex M line, is being invested in heavily by many microcontroller manufacturers, and it represents a large and growing portion of microcontroller sales.

To extend the battery life of sensor nodes, we seek to reduce the energy per operation of sampling and preprocessing our input sensor data. To do so, we put the CPU and analog components into sleep mode for as much time as our application will allow.

However, sleeping the analog components between sample acquisitions poses a problem for real time systems because the process of putting each component to sleep and waking it up consumes a significant amount of time and energy. We have found that for many microcontrollers, waking the ADC from sleep mode takes roughly the same amount of time and energy as acquiring a single sample. When sampling medium and high-frequency signals, the amount of time and energy consumed in powering down the analog subsystem between samples contributes significantly to the overall energy consumption of the analog to digital conversion process [27].

But sleep-wake overhead is not the only inefficiency in ADC sampling. In applications that need to sample multiple analog input channels at high frequency, the overhead of switching between channels is often a dominant contributor to the overall sampling period. This can significantly reduce the maximum sampling frequency of the ADC. Multi-channel ADCs, which are common in mid-range microcontrollers, often use an analog multiplexer to direct signals from the input channels to the conversion circuitry (see Figure 1) rather than having an independent conversion circuitry for each channel. On many microcontrollers, the analog input mux is slow to switch input channels. Applications that need to sample many analog input channels could use PANDA to reduce the overhead of input channel switching by sampling each input channel in bursts before switching to the next input channel. This would



**Figure 1: Schematic of a typical multi-channel ADC circuit. A single data conversion circuit performs digitization of all analog input data. An analog multiplexor selects a single analog input for conversion. Changing the selected analog input is a long-latency operation that reduces the system performance.**

reduce the overhead of switching between analog input channels, and it can significantly increase the maximum sampling frequency we can achieve. Table 1 gives an overview of ADC performance characteristics for a few popular microcontrollers. The rightmost column of the table shows the latency of switching ADC input channels between samples relative to sampling the same channel sequentially.

With PANDA, we seek to reduce energy and time wasted by repeatedly sleeping and waking a sensor node between samples. Figure 2 shows a comparison of (a) conventional sampling performed at uniform time intervals and (b) and (c) the nonuniform paradigm used by PANDA. Numbered circles represent samples collected by the microcontroller. Rather than waking the CPU and ADC for a single conversion at uniform time intervals, we propose collecting samples in bursts. After waking the microcontroller from sleep, several samples will be collected from the ADC in rapid succession before returning to the sleep state. Sampling in this way reduces the number of times we need to transition the CPU in and out of its low-power sleep state.

In Figure 2 (a), using uniform sampling, we transition the CPU in and out of sleep a total of six times, using overhead energy each time. In Figure 2 (b), we collect samples in bursts of two. Here, we collect the same number of samples in the same amount of time, but we only have to transition the CPU in and out of sleep three times. This reduces the amount of energy spent on overhead by half. In Figure 2 (c), we further reduce the energy overhead by collecting samples in bursts of three, transitioning in and out of sleep only twice and saving more energy still.

Unlike other burst-mode data acquisition techniques, PANDA provides meaningful frequency domain analysis tools including frequency spectrum analysis, linear filtering, and signal reconstruction. The availability of frequency domain signal processing techniques makes PANDA a practical tool for real-time signal processing applications like energy metering, biometric (ECG, EEG) monitoring, and so forth.

Intuitively, we think of burst-mode sampling as capturing information about a signal’s value and its first few derivatives. In a short burst of four samples, for example, we can record the signal’s value and its first three derivatives. Knowledge of the signal’s first few derivatives at a moment in time gives us information about how

the signal will evolve in the near future. This makes it possible for us to skip samples after the burst, which saves energy.

Unfortunately, direct Fourier analysis techniques—FFT, DFT, etc.—are not useful for interpolation of nonuniformly sampled signals because the derivatives of a signal are linearly dependent. Fourier techniques are successful at interpolating signals because the basis of sinusoids they use are orthogonal. To solve this problem, we use other methods of interpolating and resampling the signal before applying Fourier methods.

Under the burst sampling paradigm, we can amortize the overhead of bringing the microcontroller out of its sleep state over more than one ADC sample, allowing us to save time and energy. However, the data produced by PANDA is not sampled at uniform intervals, and therefore cannot be fed directly into conventional signal processing algorithms such as the FFT.

To deal with this problem, we piggyback off of work that has been done in the signal processing community in the area of nonuniform sampling and reconstruction. We adapt the existing work, which was initially developed for use in RADAR systems, to work on microcontrollers. We then evaluate the timing implications both for ADC sampling and data processing, showing that we can achieve **more than 17% reduction** in sampling energy compared to uniform sample spacing, even when accounting for the additional processing overhead required to compute the frequency spectrum. To our knowledge, PANDA is the first algorithm to adapt nonuniform DFT techniques for online use in a resource-constrained runtime environment.

The key contributions of this work are as follows:

- We adapt an existing nonuniform DFT algorithm to work online in a resource-constrained runtime environment.
- We compare the energy used by PANDA to conventional uniformly spaced sampling, demonstrating that burst-mode sampling can **reduce the energy per sampling operation by more than 17%**, including additional computation overhead.
- We describe the limitations of frequency domain analysis on burst mode sampled data that we encountered in resource constrained real-time environments.

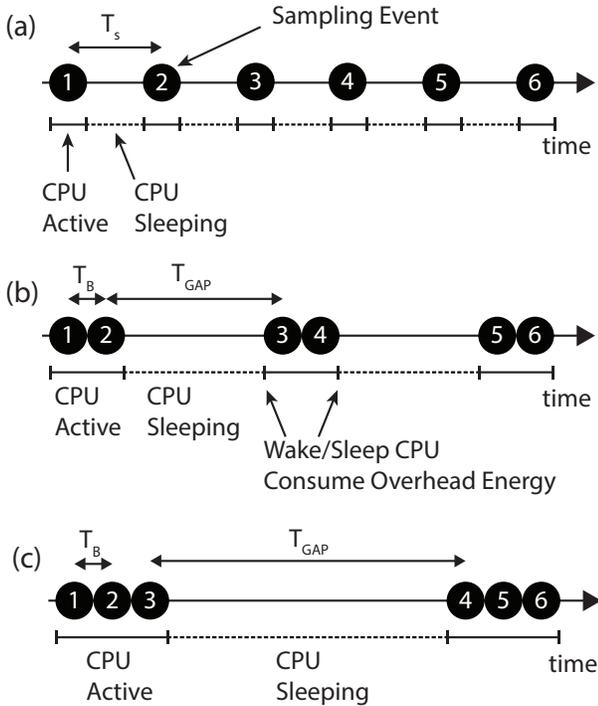
## 2 BACKGROUND

The energy required to perform analog to digital conversions in commodity microcontrollers varies by architecture, but it can account for a significant portion of energy budget. For low-power applications, the MCU will enter a sleep state, powering down unnecessary components whenever possible, in order to conserve energy. Most microcontrollers support several different levels of sleep, where “deeper” sleep modes correspond to powering down more functions or peripherals of the microcontroller. When entering deep sleep states, the sensor node must typically reconfigure clock generation and distribution circuitry such as the microcontroller’s phase-locked loop (PLL) and clock tree, a process which consumes time and energy. So each time the sensor node enters or returns from a sleep state, it must spend time and energy to do so.

For nodes like temperature sensors that take data samples at extremely low frequency, the energy consumed by sleeping and

MCU	Number of Channels	Same Channel	Alternate Channels	Slowdown
Atmel AVR XMEGA	16	$0.25\mu s$	$3.5\mu s$	14x
Freescale Kenetis K50-72MHz	2	$1.25\mu s$	$10\mu s$	8x
Freescale ColdFire	24	$5\mu s$	$15.95\mu s$	3.19x
TI ARM Cortex M4	24	$1\mu s$	$3\mu s$	3x

**Table 1: A comparison of ADC sampling time among several popular commodity microcontrollers. Same channel is the amount of time it takes to sample the same ADC channel sequentially. Alternate channel is the amount of time it takes to sample different channels sequentially. Slowdown is the relative performance penalty of alternating channels compared to sampling the same channel sequentially.**



**Figure 2: A comparison of (a) conventional sampling performed at uniform time intervals to (b) and (c) nonuniform burst sampling used by PANDA. In this example, we amortize the overhead of sleeping and waking the microcontroller over three samples by collecting data in bursts.**

waking the CPU and peripherals is insignificant. In higher frequency applications, on the other hand, which require sampling real-time signals at rates of more than few hundred Hertz, the overhead of repetitively sleeping and waking the microcontroller represents a large fraction of the node’s overall energy consumption. For extremely low-power applications, this overhead can be comparable to the amount of energy used by the communications interface [27].

Figure 3(a) shows the power consumed by a 32-bit ColdFire microcontroller’s ADC and CPU during a single ADC sample. After waking from its sleep state, the MCU spends a relatively short amount of time—roughly  $14\mu s$ —sampling its ADC. Once the ADC has been sampled, the CPU and ADC return to a sleep state to

conserve power, a process that takes more than  $100\mu s$ . For a single sample, most of the energy is consumed during the power down phase, while components in the MCU return to their sleep state.

In contrast, Figure 3(b) shows the power consumed by the ColdFire during a burst of four sample acquisitions. In burst mode, the overhead of powering down the CPU and ADC can be amortized over several sample acquisitions. This is the key observation that allows PANDA to conserve power.

Figures 3 (c) and (d) show the power consumed by an 8-bit Atmel microcontroller in 1- and 4-sample bursts. Although it is clearly far more efficient in terms of its sleep overhead power consumption, the Atmel microcontroller exhibits similar tendencies.

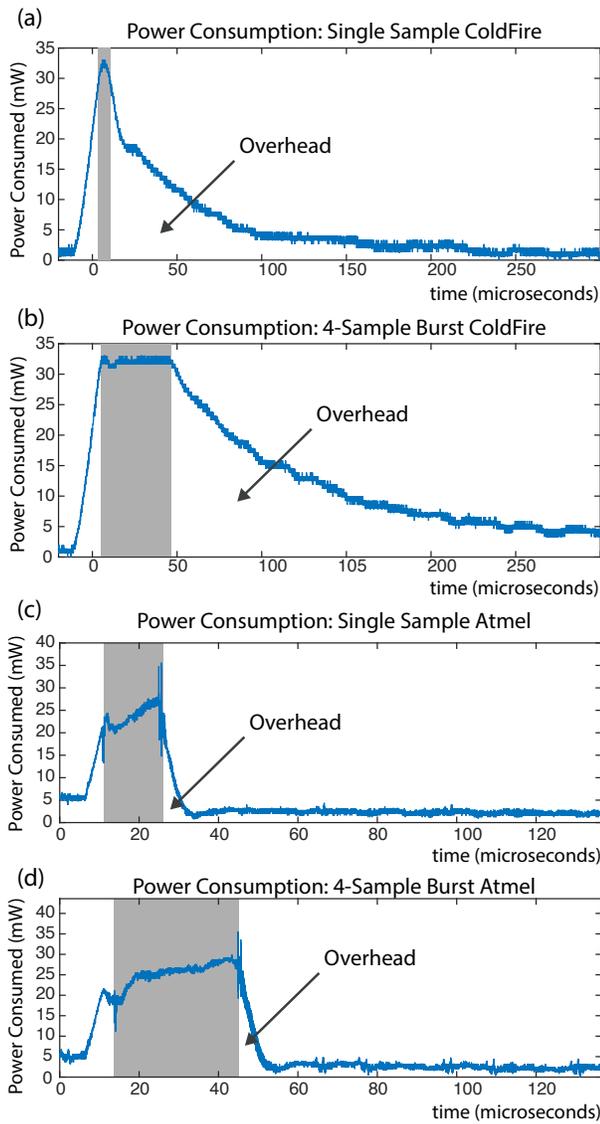
By sampling our signal of interest in in high-frequency bursts followed by extended periods of deep sleep, we can effectively gather the same information about our target signal while conserving energy by reducing the overhead required to repetitively wake and sleep the MCU. Figure 4 shows how the overhead energy per sample decreases for increasing sample burst lengths. This approach has been used in many systems [9], but until now, only rudimentary signal processing techniques could be applied<sup>1</sup>.

Once the data is collected, we want to be able to analyze it using frequency-domain techniques. For example, we may wish to apply a linear filter to the data and then resample it uniformly in time to be analyzed further. To do so, we apply a matrix transform operator, much the same as the matrix formulation of the conventional Discrete Fourier Transform.

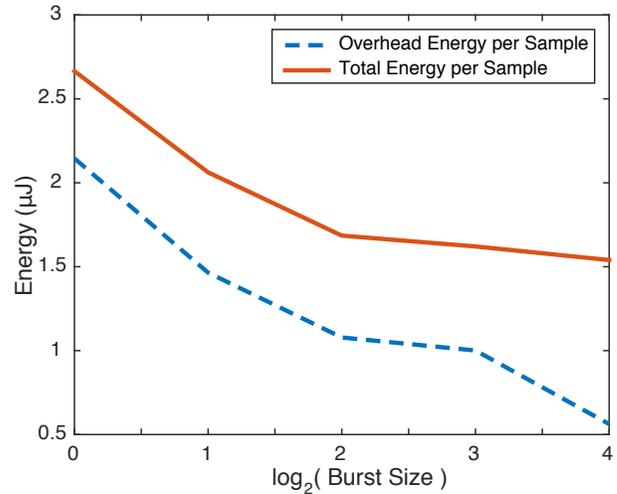
### 2.1 Related Work

There has been a lot of study on reconstruction of signals from nonuniform samples. Analysis of nonuniformly sampled signals was first investigated by Yen [35] in 1956. Early work in the area was concerned with analyzing RADAR return signals. Since that time, techniques have been developed to analyze nonuniformly sampled signals in astronomy [22], medicine [21], image processing [5, 25, 34], and others. Laguna and Moody studied the power spectral density of unevenly sampled data using least square analysis and modeled the uneven sampling as uniform sampling plus a stationary random deviation. They applied this technique to analyze the Heart Rate Signals [21]. Theoretical contributions have been made to improve the quality of frequency spectrum estimates and reconstruction techniques [24]. Jenq talks about perfect reconstruction of digital spectrum from the nonuniformly sampled signals

<sup>1</sup>For example, burst mode sampling is commonly used to capture the value of a DC signal in additive white Gaussian noise by computing the mean value of many successive samples before returning the node to sleep



**Figure 3: Power consumed by the CPU and ADC to acquire (a) a single ADC sample and (b) a burst of four samples in succession. After waking from sleep, the CPU and ADC power up, acquire samples, and return to sleep. The area highlighted in gray shows the time during which the microcontroller is doing useful work. For the rest of the time—a short powerup beforehand and a long powerdown sequence afterward—represent overhead during which time the MCU is consuming power but not performing any useful work. By amortizing this overhead over bursts of multiple ADC samples, PANDA reduces the overall energy consumption of ADC sampling.**



**Figure 4: Total energy and energy overhead per sample for different lengths of sample burst. As bursts get longer, we can amortize the overhead of waking and sleeping the microcontroller over more samples. In this application, we see diminishing returns above 4 samples per burst.**

where the timing offsets of each sampling instance are known and have a periodic structure [17].

Butzer and Hinsien also offer some interesting insight reconstruction of bounded signals from pseudo-periodic, irregularly spaced samples [4]. Moinch and Boche’s work on the non-equidistant sampling for bounded signals considers sampling patterns that are made of zeros sine-type functions and analyze the local and global convergence behavior of the sampling series. They also discuss the influence of oversampling on the global approximation behavior and the convergence speed of the sampling series [28]. However, these techniques have generally been applied only when no other course of action is practical. That is, people do not generally sample signals nonuniformly when other options are available because less is known about the mathematical properties of spectral analysis and reconstruction of nonuniformly sampled signals. Nontraditional data acquisition techniques such as compressive sensing [10] have also been applied in sensor networks to reduce power consumption. [33] used this technique for scheduling the measurement for the estimation of soil moisture. The Sparse Fast Fourier Transform [14, 15] is another technique for reducing the computational overhead of computing the most prominent harmonics of a signal.

Energy profiling of embedded systems has been widely studied topic. Zhai studied nonuniform sampling of sensors on smartphones [36], using similar techniques to decrease energy consumption by clustering sensor sampling. Tiwari et al. [31] talks about power modeling of embedded processor to characterize power budget of different software running on a processor. Cycle accurate energy estimation technique has been presented in [32]. Characterizing energy split up in the overhead is out of the scope of this work and can be done as a future work using the techniques mentioned above. Our work focuses on finding the overall energy overhead of each sample and thus we went ahead measured energy used by the

Algorithm	Computation Complexity	Stability
Marvasti [25]	$O(N^2)$	Low
Dutt [11]	$O(N \log N + N \log(1/\epsilon))$	High
Andersen [1]	$O(N^2)$	Moderate
Greengard [13]	$O(N \log N + N \log(1/\epsilon))$	High
Press [12, 29]	$O(N \log(N))$	High

**Table 2: A comparison of nonuniform DFT algorithms. Here,  $N$  is the size of the input dataset.**

entire board. Since this is a straight forward method, we have not come up with any estimation model.

### 3 ALGORITHM OVERVIEW

Broadly speaking, there are two classes of algorithms for computing the nonuniform DFT:

- (1) **Direct computation methods** use a matrix-vector multiplication, similar to the matrix form of the standard DFT:

$$f = As$$

where  $s$  is the signal of interest (a vector) sampled at nonuniform time intervals,  $A$  is a matrix whose rows are complex sinusoids sampled at the same times as  $s$ , and  $f$  is the computed frequency spectrum, also a vector. The advantage of this method is that it is a convenient and intuitive representation of the frequency spectrum. Its main disadvantage is that it tends to amplify noise in the input signal. Since the rows (and columns) of  $A$  are not orthogonal, frequency components from different harmonics of the input signal  $s$  can be confused. This effect becomes more pronounced as the sampling nonuniformity increases. For that reason, it is not suitable for burst mode sampling, which relies on highly nonuniform sampling patterns to conserve energy in the ADC.

- (2) **Interpolation and resampling methods** use some form of interpolation—for example, polynomial or spline interpolation—of the input data to construct a continuous-time model of the signal of interest. This continuous time model is then resampled at uniform intervals, and the uniformly sampled signal is fed in to a standard FFT. The interpolation and resampling approach has the advantage of being more stable for highly nonuniform sample times. A disadvantage of this approach is that it is less computationally efficient because it must solve two matrix-vector problems, compared to one in the direct method.

The fastest and most stable NUFFT algorithm we are aware of in the literature is one published by Greengard and Lee [13] for processing RADAR signals. Their algorithm, designed to run on PCs, uses a technique they call *fast Gaussian gridding* to interpolate between nonuniformly spaced samples. The idea is to replace each sample—conventionally represented as a Dirac delta function—with a Gaussian bump. The continuous time signal is then represented as an expansion Gaussians:

$$f(t) = \sum_{n=0}^{N-1} a_n e^{-(t-t_n)^2/\sigma_n^2} \quad (1)$$

In the Gaussian gridding approach, Greengard and Lee find the  $a_n$  and  $\sigma_n$  that best match the samples and then interpolate uniformly spaced samples using the basis functions. Normally, computation of those parameters would require solving an  $N \times N$  system of equation, using  $O(N^3)$  operations. Fast Gaussian gridding reduces that time by making the simplifying assumption that each Gaussian bump only affects signal values in its immediate neighborhood. This optimization produces a banded matrix which is approximately zero everywhere except on the diagonal and one or a few off-diagonal rows. Fast Gaussian gridding achieves higher performance by only considering a small subset of the basis functions for each point in the interpolation.

We have modified the Greengard and Lee algorithm for use in real-time applications in resource-constrained environments:

- We have separated the initialization phase of the algorithm into its own function, called once at startup.
- We pre-compute the values of the exponentials inside the sum of Equation 1 at compile time. This is possible because, in burst-mode sampling, we know the times  $t_n$  at which samples will be taken.

These optimizations reduce the runtime of fast Gaussian gridding by roughly 40%. At runtime, we only need to solve the banded matrix-vector system of equations to interpolate between sample bursts followed by a conventional FFT.

## 4 EVALUATION

### 4.1 Experiment Setup

In this section, we study how burst-mode sampling affects energy consumption in sensor nodes. Compared to conventional uniform sampling, burst-mode sampling in PANDA changes the way we collect and process sensor data. We consider energy used by data acquisition and preprocessing, roughly dividing the tasks into:

- (1) **Sampling energy** required to operate the sensors and analog to digital conversion hardware.
- (2) **Computation energy** required to compute the frequency spectrum of the sampled data.

For both operations, we compare the amount of energy used in conventional sampling to that used by PANDA. We then discuss how PANDA can be applied in several example applications. For each application, we will compare the energy spent sampling the waveform with PANDA to the energy spent with conventional sampling. We will also compare the amount of CPU time it takes to preprocess the data with PANDA to the amount of time spent with the FFT.

The applications we study in this work are:

#### Electric Power Meter

In this application, our goal is to measure the electric power consumed by some household appliance in real time. We use a current transformer to measure AC electric current and a resistor to measure AC voltage. An 8-bit Atmel ATmega328 microcontroller

samples the analog output of the sensors with its analog to digital converter and computes the power by multiplying the 60 Hz components of the voltage and current waveforms together.

$$P_{60} = V_{60}I_{60}$$

Between samples, the microcontroller shuts down the analog signal conditioning circuitry and powers down the microcontroller and peripherals. The microcontroller offers various sleep modes and we chose SLEEP\_MODE\_PWR\_DOWN and it was operated at maximum voltage of 5 volts. This mode shuts down all the generated clocks and allows operations of asynchronously mode only. The controller wakes up when watchdog times out, configured for a 4 seconds timeout and samples the input the signal. For the purpose of finding energy overhead per sample, we sampled 1, 2, 4, 8, and 16 samples every time the controller wakes up. We recorded the voltage vs time measurements across the resistor and computed the current through the circuit. Also, the rise and fall of voltage level at the GPIO port was also recorded and was used as the reference to mark the energy overhead. The recorded waveform was exported and used matlab to compute the area under the curve. We use frequency domain analysis (FFT) to compute the 60 Hz components of the waveforms of interest.

### Water Flow Sensor

The water flow sensor tracks the volume of water passing through a piece of industrial water treatment equipment. The water sensor hardware outputs a square voltage waveform with a frequency proportional to the water flow rate. A 32-bit ARM microcontroller records the frequency generated by the flow sensor and converts it to flow rate.

To compress the data, our microcontroller transforms the computed flow rate into the frequency domain and discards low-amplitude harmonics.

Figure 5 shows the average amount of energy used per sample in different burst lengths, with bursts of length 1 corresponding to uniform sampling. Clearly, burst mode sampling reduces the energy required to sample the signal of interest in both applications. We will study the energy requirements of computing the frequency spectrum from burst mode samples in later sections.

Figure 6 shows a breakdown of the energy used in sampling for the power meter application using various sample burst lengths. We have divided the energy consumed into (1) *sampling energy* and (2) *overhead energy*. Sampling energy is the actual energy required to run the ADC, CPU, and analog circuitry to acquire the samples. This is the energy blocked out in gray in Figure 3. Overhead energy is the energy required to power up and power down the microcontroller before and after gathering samples.

This data was gathered by setting one of the microcontroller’s GPIO lines to a 1 during the actual sampling process and setting it to a 0 just before sleep. We then gathered power consumption traces with an oscilloscope and integrated the traces numerically.

The overhead energy remains roughly constant for all burst lengths we studied because the energy required to sleep and wake the microcontroller is a constant, independent of the amount of

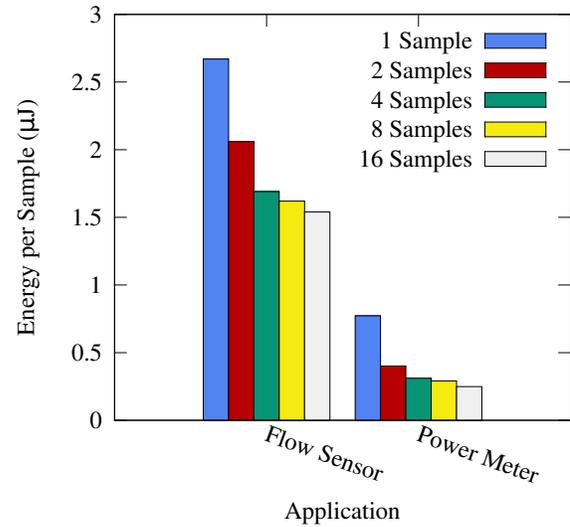


Figure 5: Comparison of energy per sample for two different application using several different burst lengths.

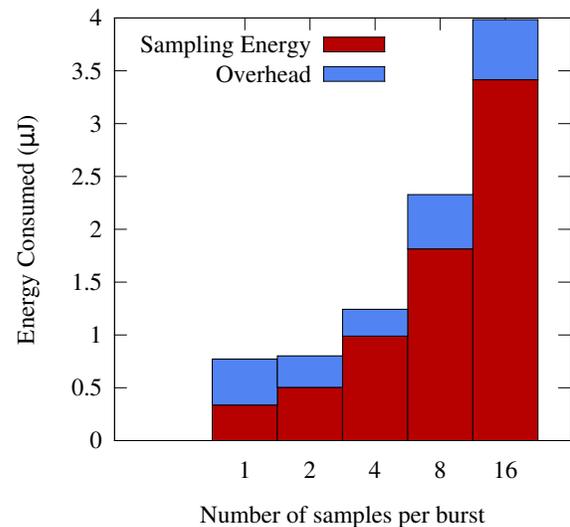


Figure 6: Energy consumed by sampling in the power meter application (AVR microcontroller) using different burst modes.

work performed. The sampling energy grows roughly linearly with the number of sample acquired.

As we increase the number of samples gathered per burst in Figure 6, we can see that the overhead energy becomes a smaller fraction of the total. Using this fact, we can amortize the energy required to sleep and wake the CPU over more samples.

## 4.2 A Simple Demonstration

Here, we will demonstrate that existing algorithms can correctly compute the frequency spectrum of nonuniformly sampled data. We also show that the computed spectrum can be fed directly into a standard inverse FFT library routine to yield a time-domain signal resampled at uniform intervals. This is important if we intend to apply standard frequency-domain analysis or filtering techniques to the nonuniformly sampled data.

Figure 7 shows a simple demonstration of how PANDA operates on a synthetic dataset. Part (a) shows the original waveform, a clipped sine wave, before any transformations have been applied. We imagine that this is the analog signal of interest that is input to our system. Its frequency spectrum  $F(\omega)$  is depicted in part (b).

Part (c) shows  $f_{AA}(t)$ , the signal after it has been passed through an analog low-pass antialias filter. This is the signal that would actually be sampled by our ADC, with high frequency components removed. The spectrum of  $f_{AA}(t)$  is shown in part (d).

Part (e) shows the samples of  $f_{AA}(t)$  that would be taken by our ADC in burst mode. In this example, we use a burst length of two samples. Each burst of samples is followed by a period of inactivity equal in length to the burst. We then compute the nonuniform DFT of the burst-mode sampled signal, shown in part (f). This should be the same as (or very close to) the frequency spectrum we would have gotten had we sampled  $f_{AA}(t)$  uniformly and computed the FFT with standard methods. Figure 7 (h) shows the error between the PANDA spectrum of  $f_{AA}(t)$  and the spectrum computed with the FFT. Stable algorithms should keep this error as small as possible.

At this stage, the data gathered using PANDA is in the same format as if we had collected samples at uniform time intervals and transformed it using the conventional FFT. We could apply standard frequency domain filtering techniques to the data in Figure 7 (f). To reconstruct the time-domain samples at uniform intervals from the filtered data, we can apply the standard inverse FFT, shown in part (h).

We assume in this example that our objective is to apply some frequency domain analysis to the nonuniformly sampled signal. If, on the other hand, all we want is to resample a nonuniformly sampled signal at uniform time intervals, there are several direct techniques for doing so [4, 23, 26].

### Requirements for Sampling Functions Nonuniformly

Marvasti et. al. provide an analysis of the requirements for nonuniformly sampled functions, which we will not repeat here [25]. They show that as long as the average sampling rate satisfies the Nyquist criterion, the continuous signal can be reconstructed without loss of information.

### Time Complexity of PANDA

In this section, we compare the amount of time (and CPU cycles) required to compute the frequency spectrum of burst-mode sampled data gathered with PANDA. The purpose of this analysis is to be sure that the energy we save with burst-mode sampling is not wasted in extra computation time on the microcontroller. There are several efficient algorithms known for computing the frequency spectrum of nonuniformly sampled data (see Table 2, the fastest of

which use  $O(N \log N + N \log(1/\epsilon))$  operations for  $N$  time-domain samples and precision of  $\epsilon$ .

Figure 8 shows runtime measurements taken on an ARM CPU of a standard FFT algorithm (KISS FFT), an in-place DFT algorithm, and PANDA's nonuniform FFT algorithm.

## 4.3 Runtime Profiling

In this section, we compare the runtime PANDA with the standard FFT, analyzing real data collected from a large deployment of water flow and quality sensors. We analyze roughly 9000 data points collected at 5-minute intervals over a time span of several months. Using the raw samples as ground truth, we created burst-mode sampled data by intermittently saving and discarding data points from the original sample database. We then use the data collection platform (which includes an ARM CPU) to process the data with the PANDA NFFT algorithm and compare the results to uniformly spaced sampling mode.

Figure 8 compares the execution time of two library DFT algorithms with PANDA. These experiments were done on a ARM microcontroller, which is a commodity MCU used in our water flow sensor application.

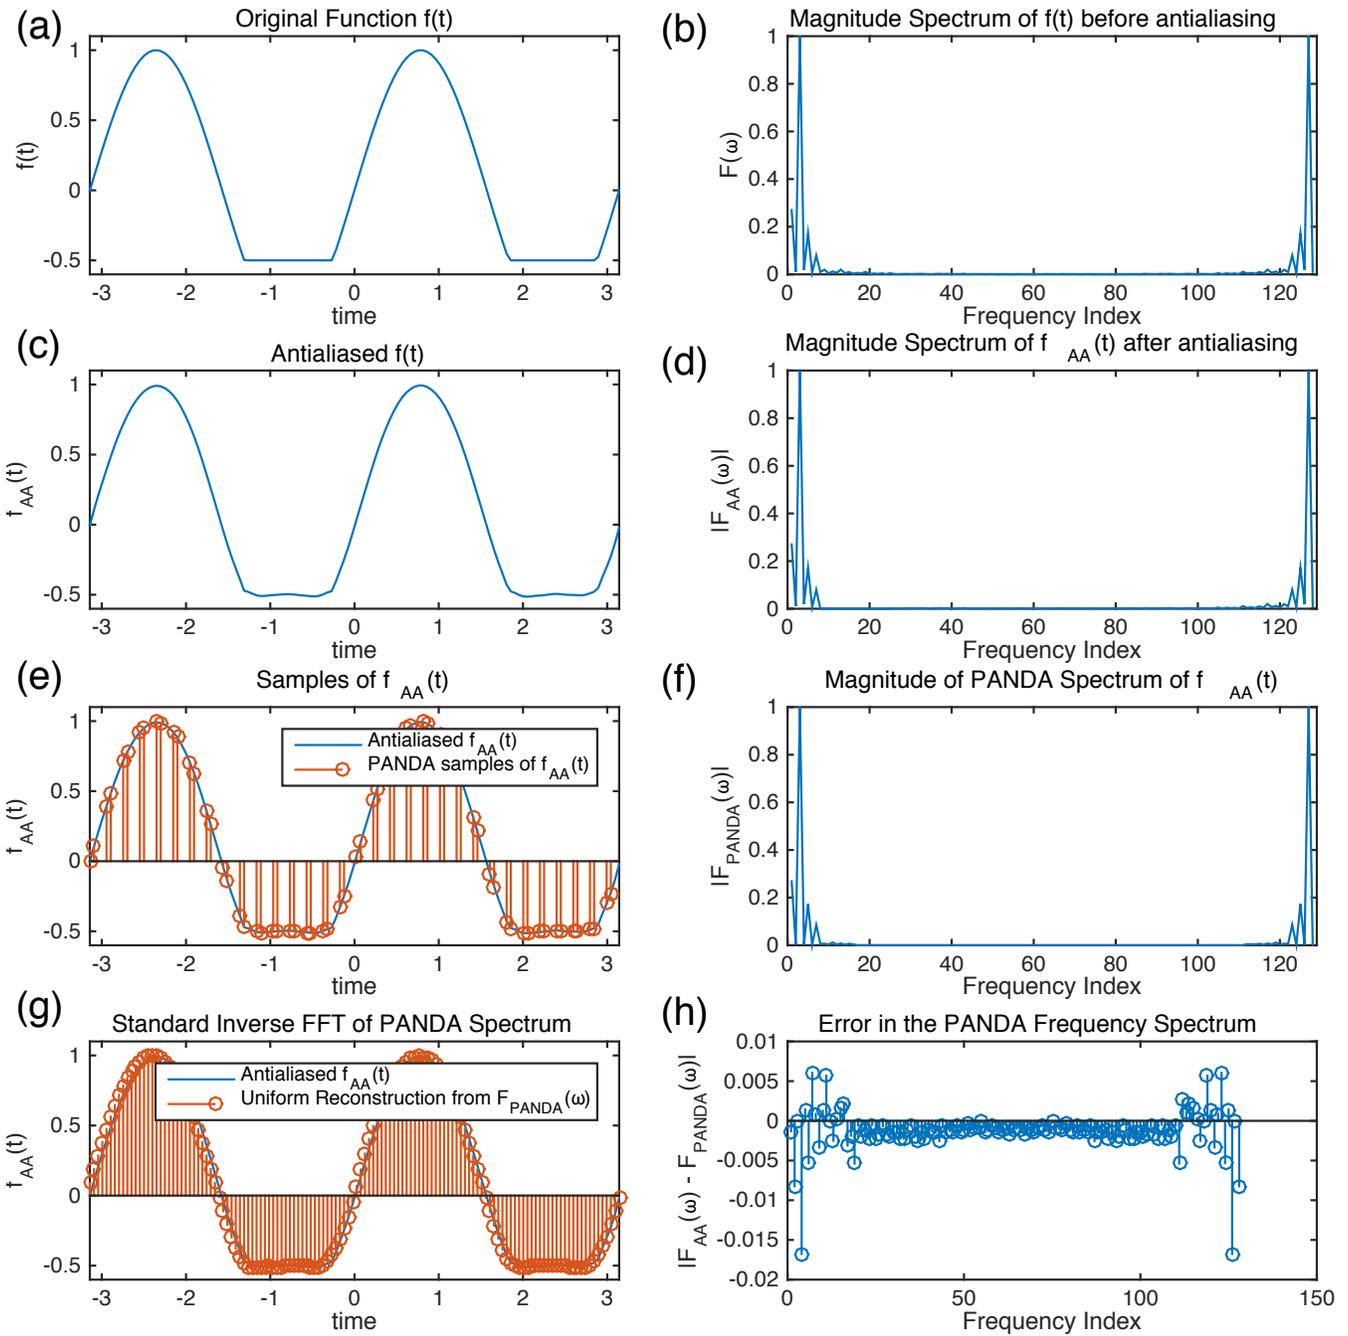
- (1) **KISS FFT** [3] is a (relatively) simple implementation of the FFT designed to run on PCs with floating point support. KISS is an out-of-place implementation of the FFT, meaning that it requires two separate buffers (one input and one output) to compute the frequency spectrum.
- (2) **In-Place FFT** [30] is an in-place DFT implementation targeted at low-end embedded systems. Since it is an in-place implementation, it requires roughly half the data storage as KISS because the input and output buffers are the same.
- (3) **PANDA NUFFT**, to compute the frequency spectrum of a nonuniformly sampled input dataset.

Surprisingly, the PANDA NUFFT runs faster on this CPU for  $2^{10}$  samples than for  $2^9$  because of the behavior of the microcontroller's data cache. We did not optimize our code for the microcontroller's cache configuration. The KISS FFT implementation takes advantage of optimizations which reduces the computational complexity to  $O(n \log(n))$ . The in-place DFT has a higher computational complexity than the FFT [6], because it cannot take advantage of FFT optimizations which require separate memory buffers to perform decimation in time. For the remainder of this paper, we will compare PANDA to the KISS FFT implementation because it can achieve a shorter runtime which will translate into less energy consumed by the CPU during the computation of the frequency spectrum.

## 4.4 An Alternative: Downsampling Without Bursts

Do we actually gain anything by sampling in bursts at lower frequencies? What would happen if, instead of sampling in bursts at lower frequencies, we just use uniform sampling at lower frequencies? We want to be sure that burst mode sampling actually gives improved resolution in the frequency domain.

Figure 9 shows how frequency resolution (on the  $y$  axis) depends on sampling intervals in our water quality dataset, with bigger numbers corresponding to worse resolution. The original dataset has samples taken at five minute intervals. To learn how frequency

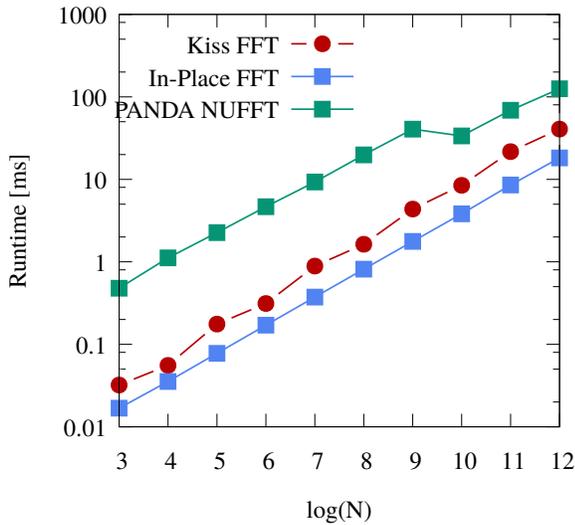


**Figure 7: A demonstration of sampling a time-domain signal at nonuniform time intervals, computing its frequency spectrum, and reconstructing the original waveform using a standard inverse FFT.**

domain resolution depends on sampling frequency, the original data was downsampled, which was possible because the original sampling frequency was much higher than the Nyquist frequency of the underlying signal.

Downsampling was done in two modes: uniform in which the dataset was downsampled uniformly in time, and burst in which

the dataset was downsampled in burst mode. Burst mode downsampling on the  $y$  axis is annotated as (number of points retained):(number of points discarded). For example, 2:4 burst downsampling means that two data points from the original dataset were retained, and the following four were discarded. Between groups, bars with the same colors correspond to samples taken at the same frequency in



**Figure 8: Runtime comparison between In-Place and Out-of-Place FFT on an ARM Microcontroller. The x-axis is the log of the dataset length, and they y-axis is the runtime of each algorithm. The jog in PANDA’s runtime is caused by L1 data cache misses for smaller sample lengths.**

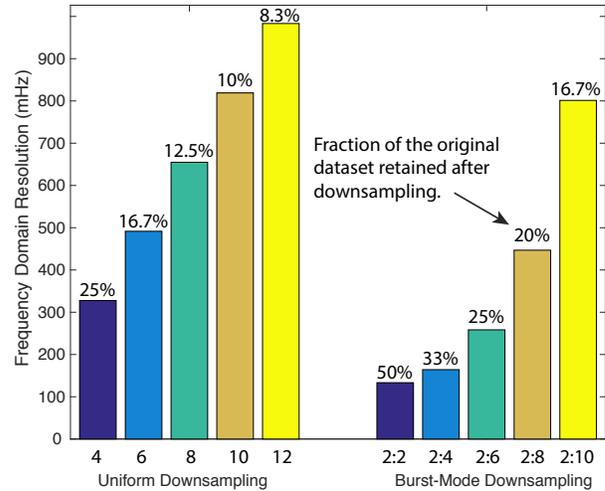
uniform sampling mode or burst sampling mode. For example, the blue bar corresponds to bursts of length two (keep two samples, throw out two samples) occurring at the same frequency as uniform downsampling by a factor of four (keep one sample, throw out three samples). Both blue bars represent sampling of the original signal at a period of  $5 \times 4 = 20$  minutes. In burst mode downsampling, we take two samples every 20 minutes, and in uniform downsampling, we take one sample every 20 minutes. We want to confirm that by retaining more data in burst mode, we actually get better resolution in the frequency domain.

As we would expect, resolution in the frequency domain gets worse (coarser) as we downsample the data more. However, in burst mode, it is clear that resolution in the frequency domain degrades more slowly, despite the fact that the time interval between bursts is the same as the corresponding time interval between uniform samples in the other bar cluster. In terms of information content in the frequency domain, there is value in taking more than one sample in a bursts. If we take more than one sample in a burst, we can extract more information about the frequency domain content

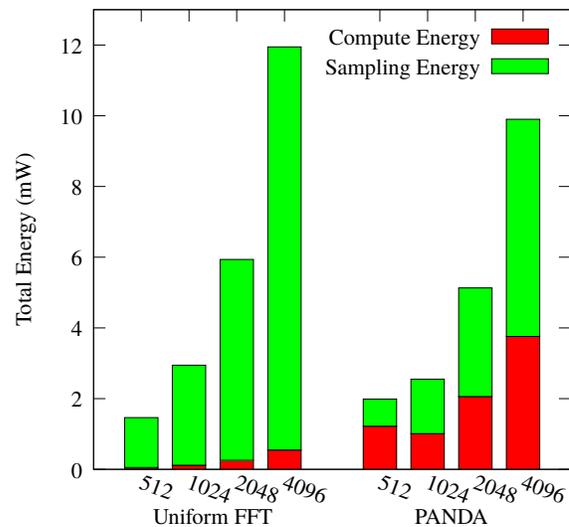
Furthermore, for comparable amounts of retained data, Figure 9 shows that we get comparable resolution in the frequency domain. For example, 2:6 burst mode sampling retains 25% of the data from the original dataset (two samples kept, six discarded). Its resolution in the frequency domain is comparable to  $4 \times$  uniform downsampling, which also retains 25% of the original data.

### 4.5 Energy Profiling

We have demonstrated that burst-mode data acquisition, when paired with the PANDA NUFFT algorithm, can achieve similar resolution in the frequency domain as uniform data acquisition.



**Figure 9: Frequency domain resolution as a function of sampling frequency for uniform downsampling and burst mode downsampling. Small is good, big is bad.**



**Figure 10: Total energy used in sampling and frequency spectrum computation for the uniform KISS FFT and the PANDA NUFFT. The x-axis is the dataset size for each algorithm, and the y-axis is the total energy used to sample and compute the frequency spectrum. PANDA outperforms the uniform FFT in terms of energy consumption for data windows larger than 1024.**

The process of sampling in burst mode clearly uses less energy than the uniform sampling alternative, but the data postprocessing takes more CPU time, as outlined in Figure 8. When we put all steps of the process together, how does the overall energy compare between the two techniques?

In Place	KISS FFT	PANDA
0.89 kbytes	1.98 kbytes	38.4 kbytes

**Table 3: Code size for each of the three algorithms we evaluated. The NUFFT algorithm used by PANDA uses roughly an order of magnitude more program memory than the other uniform FFT implementations.**

Using data gathered from our deployment of water quality sensors, we compared data acquisition and postprocessing with the PANDA algorithm and standard uniform sampling using various dataset sizes. Runtime profiling and energy measurements were done on the water quality sensors using historical data that had been gathered over the course of several months. Results of total energy consumption for sampling and frequency spectrum computation are shown in Figure 10.

In the burst-mode sampling used by PANDA, computation energy is a larger component of total energy consumption than in uniform sampling mode. In our experiments with the water sensing platform, PANDA's burst mode sampling outperform uniform sampling for datasets larger than 1024 samples. We get a 13% reduction in overall energy per sample for data windows of 4096.

#### 4.6 Memory Requirements

We compiled each of the three algorithms for an ARM microcontroller using the Thumb-2 instruction set and the `-Os` compiler flag to optimize for size. The code space used by each algorithm is shown in Table 3. None of the algorithms uses a significant amount of data memory—the only data memory used is for storing the input and output buffers. PANDA uses roughly an order of magnitude more code space than the uniform FFT implementations. This could be a drawback for low-end microcontrollers with constraints on code memory. However, with rapid advancements in integration on new microcontrollers, devices with code memory sizes in the range of a few megabytes are now commonly available, making PANDA's NUFFT implementation feasible.

#### 4.7 Limitations of PANDA

In the right execution environment—one with enough memory and CPU throughput to support fast Gaussian gridding—PANDA can reduce the average energy required to sample a signal and compute its frequency spectrum. On low-end microcontrollers, such as the Atmel device used in our power meter application, we were not able to compile the PANDA code because of code size constraints. We found that on other 8-bit platforms that the fast Gaussian gridding algorithm performs poorly because of lack of instruction set support. In particular, PANDA requires native support for integer division and barrel shift, which are not usually available on low-end 8-bit platforms.

In PANDA, we assume that the microcontroller's CPU must be awake in order for the ADC to take a sample. Some more sophisticated microcontrollers have the ability to trigger ADC conversions without intervention from the CPU and transfer the recorded value directly to memory. This feature is only available on highly-integrated devices, which are unlikely to be used in low-power

applications. PANDA is targeted at mid-range devices that are commonly used in low-power applications.

## 5 CONCLUSION

In the previous section, we demonstrated the utility of PANDA by sampling a signal in burst mode, computing its frequency spectrum, and computing the inverse transform to produce a uniformly resampled signal in the time domain. Using our method, we can save energy by sampling nonuniformly in the time domain while still applying standard frequency domain analysis and filtering techniques.

In our evaluation, we found that a large portion of the energy used by a sensor is consumed in the process of sampling. For this reason, we conclude that there is significant opportunity to reduce the energy consumed in the data acquisition process. However, in order to gather samples in burst mode, we need to use an alternative method of preprocessing the samples to compute the frequency spectrum. PANDA's approach of burst-mode sampling, therefore, only makes sense if the energy used in preprocessing the burst-mode data does not overwhelm the savings in the data acquisition phase. We found that PANDA can have benefits for common applications with average sized data sets.

We analyzed the energy consumed by PANDA's nonuniform DFT algorithm running on a real sensor application using real sensor data. We found that its runtime (and therefore its energy consumption) is higher than that of the conventional FFT for datasets of less than 1024 samples. However, for larger data sets, we demonstrated that PANDA can reduce the energy required to sample a signal and compute its frequency spectrum by as much as 17%.

## 6 ACKNOWLEDGMENTS

We would like to acknowledge the anonymous reviewers and our shepherd Dr. Prashant Shenoy for his detailed comments and guidance. All authors are supported in part by the following grants from the US National Science Foundation: CNS-1345293, CNS-14055667, CNS-1525586, CNS-1555426, CNS-1629833, CNS-1647152, CNS-1719336.

## REFERENCES

- [1] Chris Anderson and Marie Dillon Dahleh. 1996. Rapid Computation of the Discrete Fourier Transform. *SIAM J. Sci. Comput.* 17, 4 (July 1996), 913–919. <https://doi.org/10.1137/0917059>
- [2] Stuart Taylor Artem Demytyev, Steve Hodges and Joshua Smith. 2013. Power Consumption Analysis of Bluetooth Low Energy, ZigBee and ANT Sensor Nodes in a Cyclic Sleep Scenario. In *International Wireless Symposium*. IEEE.
- [3] Mark Borgerding. KISS FFT. (????). <http://kissfft.sourceforge.net>.
- [4] Paul Leo Butzer and G Hinsen. 1989. Reconstruction of bounded signals from pseudo-periodic, irregularly spaced samples. *Signal processing* 17, 1 (1989), 1–17.
- [5] D. Chen and J. Allebach. 1987. Analysis of error in reconstruction of two-dimensional signals from irregularly spaced samples. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 35, 2 (Feb 1987), 173–180. <https://doi.org/10.1109/TASSP.1987.1165115>
- [6] W.T. Cochran, James W. Cooley, D.L. Favon, H.D. Helms, R.A. Kaenel, W.W. Lang, Jr. Maling, G.C., D.E. Nelson, C.M. Rader, and Peter D. Welch. 1967. What is the fast Fourier transform? *Proc. IEEE* 55, 10 (Oct 1967), 1664–1674. <https://doi.org/10.1109/PROC.1967.5957>
- [7] Gabe Cohn, Erich P. Stuntebeck, Jagdish Nayayan Pandey, Brian P. Otis, Gregory D. Abowd, and Shwetak N. Patel. 2010. SNUPI: sensor nodes utilizing powerline infrastructure.. In *UbiComp (ACM International Conference Proceeding Series)*, Jakob E. Bardram, Marc Langheinrich, Khai N. Truong, and Paddy Nixon (Eds.). ACM, 159–168. <http://dblp.uni-trier.de/db/conf/huc/ubicomp2010.html#CohnSPOAP10>

- [8] CrossBow. MICA2 Wireless Measurement System. (????). <http://www.eol.ucar.edu/isf/facilities/isa/internal/CrossBow/DataSheets/mica2.pdf>.
- [9] Samuel DeBruin, Bradford Campbell, and Prabal Dutta. 2013. Monjolo: An Energy-harvesting Energy Meter Architecture. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys '13)*. ACM, New York, NY, USA, Article 18, 14 pages. <https://doi.org/10.1145/2517351.2517363>
- [10] D.L. Donoho. 2006. Compressed sensing. *Information Theory, IEEE Transactions on* 52, 4 (April 2006), 1289–1306. <https://doi.org/10.1109/TIT.2006.871582>
- [11] A. Dutt and V. Rokhlin. 1995. Fast Fourier Transforms for Nonequispaced Data. *Applied and Computational Harmonic Analysis* 2, 1 (1995), 85 – 100. <https://doi.org/10.1006/acha.1995.1007>
- [12] Brian P Flannery, Wiliam H Press, Saul A Teukolsky, and William Vetterling. 1992. Numerical recipes in C. *Press Syndicate of the University of Cambridge, New York* 24 (1992).
- [13] Leslie Greengard and June-Yub Lee. 2004. Accelerating the nonuniform fast Fourier transform. *SIAM review* 46, 3 (2004), 443–454.
- [14] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. 2012. Nearly Optimal Sparse Fourier Transform. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC '12)*. ACM, New York, NY, USA, 563–578. <https://doi.org/10.1145/2213977.2214029>
- [15] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. 2012. Simple and Practical Algorithm for Sparse Fourier Transform. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*. SIAM, 1183–1194. <http://dl.acm.org/citation.cfm?id=2095116.2095209>
- [16] Josiah Hester, Travis Peters, Tianlong Yun, Ronald Peterson, Joseph Skinner, Bhargav Golla, Kevin Storer, Steven Hearndon, Kevin Freeman, Sarah Lord, Ryan Halter, David Kotz, and Jacob Sorber. 2016. Amulet: An Energy-Efficient, Multi-Application Wearable Platform. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM (SenSys '16)*. ACM, New York, NY, USA, 216–229. <https://doi.org/10.1145/2994551.2994554>
- [17] Y-C Jenq. 1997. Perfect reconstruction of digital spectrum from nonuniformly sampled signals. *Instrumentation and Measurement, IEEE Transactions on* 46, 3 (1997), 649–652.
- [18] Neil Klingensmith, Joseph Bomber, and Suman Banerjee. 2014. Hot, Cold and in Between: Enabling Fine-grained Environmental Control in Homes for Efficiency and Comfort. In *Proceedings of the 5th International Conference on Future Energy Systems (e-Energy '14)*. ACM, New York, NY, USA, 123–132. <https://doi.org/10.1145/2602044.2602049>
- [19] Neil Klingensmith, Anantharaghavan Sridhar, Zachary LaVallee, and Suman Banerjee. 2015. Water or Slime? A Platform for Automating Water Treatment Systems: Poster Abstract. In *Proceedings of the 2nd ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys '15)*. ACM, New York, NY, USA.
- [20] Neil Klingensmith, Dale Willis, and Suman Banerjee. 2013. A Distributed Energy Monitoring and Analytics Platform and Its Use Cases. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys '13)*. ACM, New York, NY, USA, Article 36, 2 pages. <https://doi.org/10.1145/2528282.2534156>
- [21] Pablo Laguna, George B Moody, and Roger G Mark. 1998. Power spectral density of unevenly sampled data by least-square analysis: performance and application to heart rate signals. *Biomedical Engineering, IEEE Transactions on* 45, 6 (1998), 698–715.
- [22] Nicholas R Lomb. 1976. Least-squares frequency analysis of unequally spaced data. *Astrophysics and space science* 39, 2 (1976), 447–462.
- [23] Evgeny Margolis and Yonina C Eldar. 2008. Nonuniform sampling of periodic bandlimited signals. *Signal Processing, IEEE Transactions on* 56, 7 (2008), 2728–2745.
- [24] Robert J Marks (Ed.). 1993. *Advanced Topics in Shannon Sampling and Interpolation Theory*. Springer, New York, NY.
- [25] Farokh Marvasti. 2001. *Nonuniform sampling: theory and practice*. Vol. 1. Springer.
- [26] Farokh Marvasti, Mostafa Analoui, and Mohsen Gamshadzahi. 1991. Recovery of signals from nonuniform samples using iterative methods. *Signal Processing, IEEE Transactions on* 39, 4 (1991), 872–878.
- [27] ST Microelectronics. TA0342: Accurate power consumption estimation for STM32L1 series of ultra-low-power microcontrollers. (????). <http://www.st.com/>.
- [28] Ullrich J Mönich and Holger Boche. 2010. Non-equidistant sampling for bounded bandlimited signals. *Signal processing* 90, 7 (2010), 2212–2218.
- [29] William H Press and George B Rybicki. 1989. Fast algorithm for spectral analysis of unevenly sampled data. *The Astrophysical Journal* 338 (1989), 277–280.
- [30] Tom Roberts. Fixed-point Fast Fourier Transform. (????). [http://web.mit.edu/ghudson/dev/nokrb/third/gst-plugins/gst/spectrum/fix\\_fft.c](http://web.mit.edu/ghudson/dev/nokrb/third/gst-plugins/gst/spectrum/fix_fft.c).
- [31] Vivek Tiwari, Sharad Malik, and Andrew Wolfe. 1994. Power analysis of embedded software: a first step towards software power minimization. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on* 2, 4 (1994), 437–445.
- [32] Narayanan Vijaykrishnan, Mahmut Kandemir, Mary Jane Irwin, Hyun Suk Kim, and Wu Ye. 2000. *Energy-driven integrated hardware-software optimizations using SimplePower*. Vol. 28. ACM.
- [33] Xiaopei Wu and Mingyan Liu. 2012. In-situ soil moisture sensing: measurement scheduling and estimation using compressive sensing. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*. ACM, 1–12.
- [34] John I Yellott. 1990. The photoreceptor mosaic as an image sampling device. In *Advances in photoreception: Proceedings of a symposium on frontiers of visual science*. 117–134.
- [35] J Yen. 1956. On nonuniform sampling of bandwidth-limited signals. *Circuit Theory, IRE Transactions on* 3, 4 (1956), 251–257.
- [36] Shuang Zhai, Liwei Guo, Xiangyu Li, and Felix Xiaozhu Lin. 2017. Decelerating Suspend and Resume in Operating Systems. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications (HotMobile '17)*. ACM, New York, NY, USA, 31–36. <https://doi.org/10.1145/3032970.3032975>