

© 2012 by Gong Zhang.

BIO-INSPIRED ERROR-TOLERANT AND ENERGY-EFFICIENT
SIGNAL PROCESSING

BY
GONG ZHANG

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2012

Urbana, Illinois

Master's Committee:

Professor Naresh R. Shanbhag, Chair
Professor Douglas L. Jones

Abstract

In many applications at the sensory edge, such as security and environmental sensing, reliable sensor nodes must operate for extended time periods on battery supplies. To meet this constraint, energy-efficient systems have been developed through different technologies. The primary and the most effective approach has been technology scaling. Another emerging technique is to operate circuits in the subthreshold region as some of the applications such as environmental sensing do not require high throughput. However, both techniques lead to large process, voltage and temperature variations and therefore jeopardize system reliability.

In order to achieve both energy savings and system reliability, we take inspirations from biological methods, such as population-coding, and apply these methods to a canonical problem of non-coherent (unknown phase) frequency estimation. Energy efficiency is achieved using low cost, overlapping band-pass filters rather than conventional non-overlapping band-pass filters. Energy savings are also achieved by operating the hardware at a voltage lower than the nominal voltage (voltage overscaling), which leads to hardware timing errors. In the presence of these hardware errors, signal statistics are generated from overlapping band-pass filters with frequency redundancy. Robust techniques, such as median estimation and algorithmic noise-tolerance, are applied to filter outputs to achieve error-tolerance. Energy/performance trade-offs are further explored by altering the supply voltage. Simulation results show that the root-mean-squared-error of the bio-inspired method can be reduced by an order of magnitude relative to that of the conventional architecture while achieving an energy consumption reduction of 78% relative to the conventional method which is under hardware-error-free operations at nominal supply voltage.

To my family, for their love and support.

Table of Contents

Chapter 1 Introduction	1
1.1 Future Computer Workloads and Perceptual Sensor Network	1
1.2 Energy-Efficient Digital Signal Processing	1
1.3 Error-Tolerant Digital Signal Processing	2
1.4 Problem Definition	5
1.5 Timing Errors and Voltage Overscaling	6
1.6 Thesis Organization	7
Chapter 2 Bio-inspired Signal Processing and System Description	8
2.1 The Redundant Sensor Network	8
2.2 The Low-Cost Sensor	10
2.3 Signal Statistics in Time Domain under Voltage Overscaling	10
2.4 The Nonoverlapping Estimator	11
2.5 The Bio-inspired Estimator	13
Chapter 3 Simulation Setup and Results Comparison	16
3.1 Simulation Setup	16
3.2 Performance Comparison	17
3.2.1 Performance comparison	17
3.2.2 Energy consumption comparison	18
Chapter 4 Conclusion	21
4.1 Summary	21
4.2 Future Work	21
References	23

Chapter 1

Introduction

1.1 Future Computer Workloads and Perceptual Sensor Network

The past few decades have witnessed dramatic change of computational workloads, which used to be CPU centric and could be based on high-end servers and personal computers. Today's computational workload features mobile devices, which are the fastest growing sector in the consumer electronics industry in the past ten years. Future workloads can be roughly classified into four categories: high performance computing tasks, complex distributed systems, personalized services and surrounding perceptual processing sensor networks. The last category is evolving towards more advanced user interfaces [1].

Such emerging applications impose great challenges because there are new design metrics needed. For example, sensor networks running on batteries need to be energy efficient while low power consumption may be one of the traditional design constraints. Furthermore, sensor networks have more stringent requirements on traditional design metrics such as reliability, as these networks often operate in a highly dynamic and noisy environment.

1.2 Energy-Efficient Digital Signal Processing

Computational platforms on the sensory edge comprise computational cores and the input/output components such as sensors and analog-to-digital converters. The computational core is usually implemented using digital circuits

and consumes energy as given by [2]

$$E=CV_{dd}^2, \quad (1.1)$$

where C is the average switching capacitance and V_{dd} is the supply voltage. This indicates that energy consumption scales down significantly with technology scaling as both C and V_{dd} scale down with transistor size until the 130 nm technology node. Beyond 130 nm technology node, the supply voltage is kept at 1.2 V to keep the leakage current down. However, new device technologies such as FinFET [3] could further reduce the supply voltage. Nevertheless, transistor feature size scaling has been the major driving force to reduce energy consumption by digital circuit operation. However, transistor feature size scaling magnifies variations caused by manufacturing process, supply voltage fluctuations, temperature hot spots on the chip, and circuit aging effects [4].

Another recent trend in digital circuit design that achieves low energy operation at the cost of sacrificing throughput is to operate the circuit in the sub-threshold regime where the gate-to-source voltage (V_{gs}) is below the threshold voltage. As the current scales exponentially with the V_{gs} in this regime, the throughput of the system is greatly reduced [5]. Hence, sub-threshold design is suitable for certain applications such as bio-medical applications in which the sampling rate is usually in the kilohertz range. Similarly to feature size scaling, sub-threshold design achieves energy efficiency while introducing more susceptibility to process, voltage and temperature (PVT) variations [6].

In summary, technology scaling and subthreshold circuit design can reduce circuit energy consumption effectively. However, these trends make digital circuits exhibit statistical rather deterministic behavior, which leads to the challenge of error-resilient or error-tolerant digital circuit design.

1.3 Error-Tolerant Digital Signal Processing

In general, redundancy can be used to detect and/or correct hardware errors, including errors caused by timing violations in digital circuits. Redundancy can be achieved by either spatial, temporal or spatio-temporal redundancy.

Spatial redundancy in hardware is achieved by adding hardware replicas.

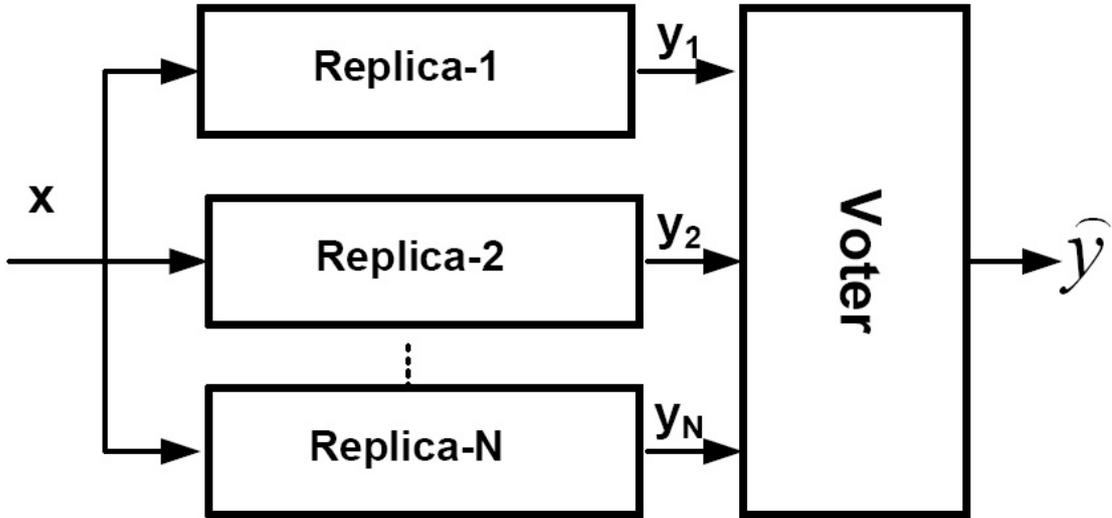


Figure 1.1: N-modular redundancy (NMR).

NMR consists of N replicas of the original system and a voter, as depicted in Figure 1.1. The most commonly used NMR is triple modular redundancy (TMR), having three replicas and a majority voter. If a single hardware error occurs in any of the three replicas, the other two units can detect and correct the error by majority voting. However, the implementation requires the voter to be error-free [7].

The check-point technique (temporal redundancy) uses re-computation or redundancy in time and divides the datapath into stages. At the end of each stage the state of the computation is stored so that computations can be rolled back and recomputed to guarantee a consistent outcome [8]. The check-point technique is an effective way to correct transient errors.

A recent work (RAZOR) can be viewed as an example of spatio-temporal redundancy. It aims to correct timing errors caused by PVT variations. Each pipeline stage incorporates two flip-flops in parallel. One has the normal clock while the other one (RAZOR flip-flop) has a clock with a delayed edge. The clock frequency and the supply voltage are set to guarantee that the RAZOR flip-flop meets the setup and hold time constraints in the worst case. At each check-point, the two latched data outputs are compared against each other. Inconsistent results would cause the computation to fall back to the previous check-point and run again while scaling the supply voltage to meet the

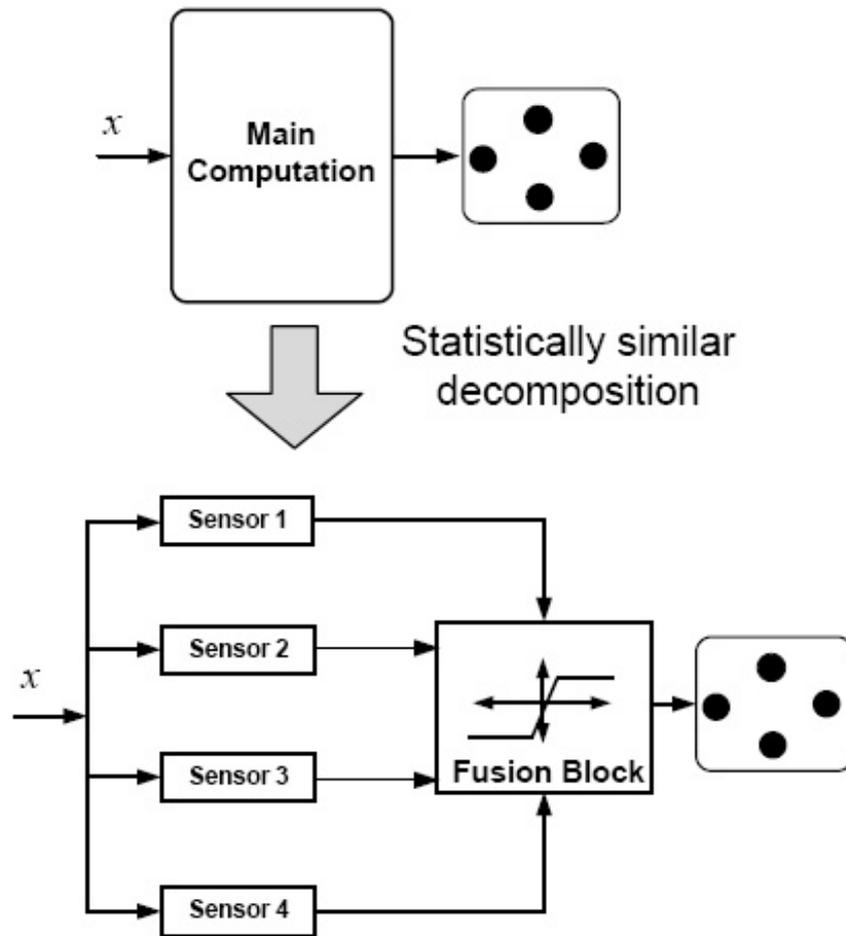


Figure 1.2: Stochastic sensor network on a chip (SSNOC) [10],[11].

timing constraints [9].

Redundancy can achieve error-tolerance at the expense of energy consumption and therefore counters the energy benefits due to transistor feature size scaling and subthreshold circuit design.

In some applications, it is possible to introduce redundancy into the system without significantly increasing the energy consumption overhead. As shown in Figure 1.2 this technique is applied to the PN-code acquisition problem. The original main architecture is decomposed into sub-systems (sensors) with similar output statistics. Results from individual sensors are fed into the fusion block, which is a median filter, to achieve error-tolerance. Redundancy is achieved with small energy consumption overhead associated with the final fusion block [10][11]. However, this technique is application specific.

In this work, bio-inspired signal processing concepts such as population-coding are explored to develop a robust and energy-efficient bio-inspired computational system.

1.4 Problem Definition

Future computing applications on the sensory edge will be characterized by very different requirements both for energy consumption and noise tolerance, with an emphasis on tasks such as sensory data processing, mining and fusion, detection and recognition, and scene and situation analysis. Neither conventional computing systems nor signal-processing algorithms perform well for these anticipated future workloads. In contrast, biological systems perform extraordinarily well in such situations. Biological systems display great robustness to variation and uncertainty, remarkable abilities to fuse data from different senses, and incredibly low energy consumption in performing these tasks.

To demonstrate bio-inspired signal processing concepts, an audio-frequency-band application is chosen. This application reflects many characteristics of sensory-edge applications. It is expected that the system design principles obtained for this application apply equally to many other types of sensory data such as acoustically steered cameras, multi-modal automated light switches, acoustic omnipresence with anyone else in the same or a similarly equipped space, and interactive toys and devices.

The problem chosen in this paper is a single-tone sinusoid frequency estimation in the audio frequency range ($2 \sim 14$ kHz). When the signal to be estimated only contains additive white Gaussian noise, the optimal estimator in term of mean-square-error (MSE) is a non-overlapping filter bank, in which the frequency estimation is given by the filter with the maximum-energy response [12].

For the problem formulated above, a certain number of samples in the time domain need to be accumulated to record the filter energy response in the time domain. In this work, 64 samples are recorded as shown in Figure 1.3. The solid line plots the filter response in time domain with an input frequency within the passband of this filter. When the input signal matches the filter, the signal-to-noise ratio is maximized and the output is considered

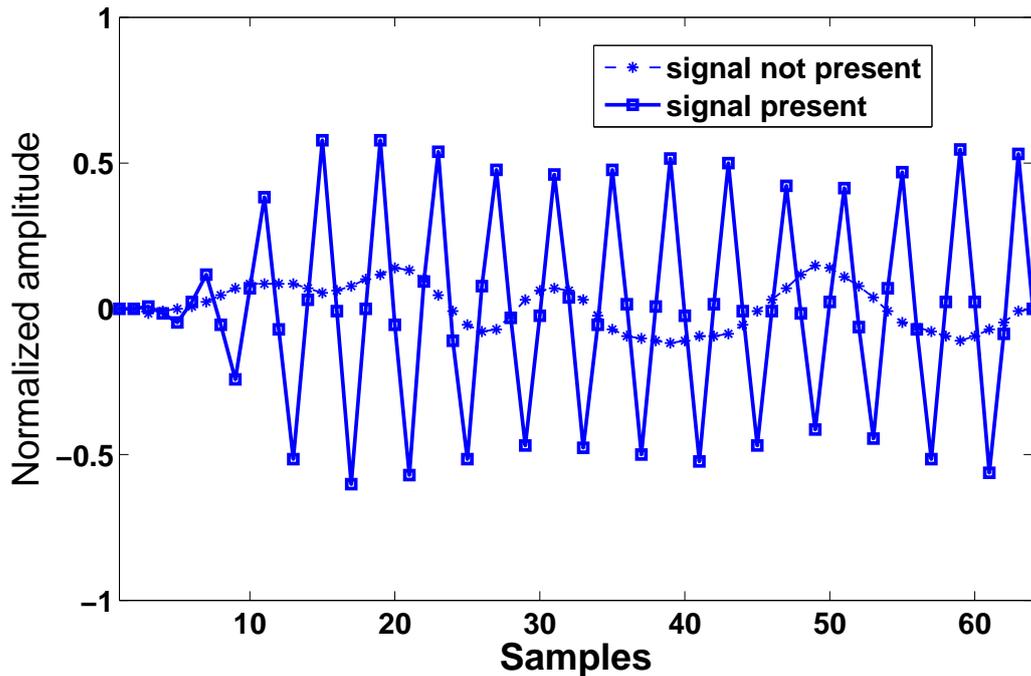


Figure 1.3: Narrowband bandpass filter response with and without the signal.

as the energy response. At the next time step, as the input signal has shifted phase and does not match the filter, the output has smaller amplitude. The dashed line shows that the filter output has a much smaller amplitude when the input frequency is outside the passband of the filter.

1.5 Timing Errors and Voltage Overscaling

This work focuses on timing errors caused by PVT variation. Timing errors are introduced into the system by deliberately lowering the supply voltage to reduce the system energy consumption. Doing so increases the critical path delay. For a given clock frequency this could cause set-up time violations and hence timing errors. Voltage overscaling is applied to the combinational logic part of the digital circuits while the registers are operated under the nominal supply voltage [6].

1.6 Thesis Organization

The remainder of the thesis is organized as follows: Chapter 2 presents bio-inspired signal processing concepts and an algorithmic noise-tolerance (ANT) filter based on signal statistics in the time domain, which are applied to the canonical problem introduced in the introduction. Furthermore, Chapter 2 formalizes the algorithms explored and the system architecture. Chapter 3 presents the simulation methodology and compares the performance in terms of root-mean-squared-error (RMSE) and energy consumption of different methods. Chapter 4 gives the conclusion and potential further research directions.

Chapter 2

Bio-inspired Signal Processing and System Description

After the specific problem of estimating a single-tone sinusoid frequency has been introduced in the last chapter, this chapter elaborates the bio-inspired signal processing concepts such as population-coding in the context of this canonical problem.

2.1 The Redundant Sensor Network

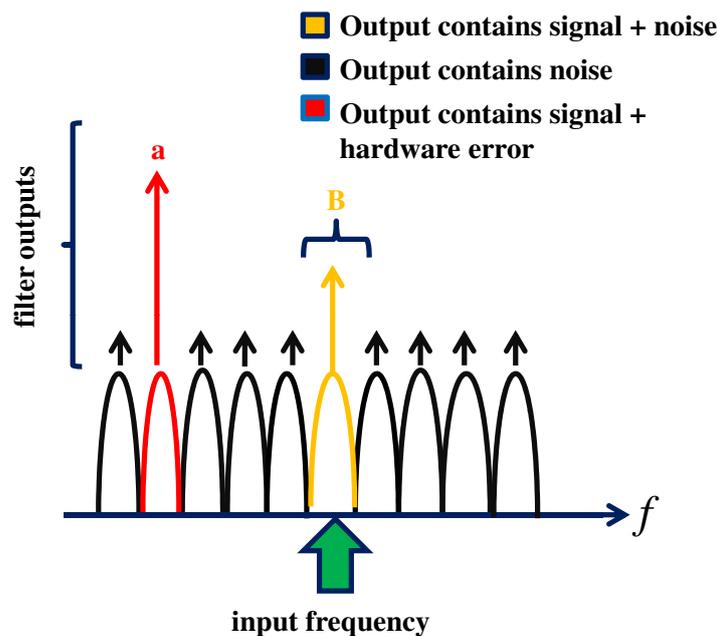


Figure 2.1: Conventional nonoverlapping filter approach. The arrow with 'a' indicates the filter output with hardware errors. The arrow with 'B' indicates the filter output with the signal in the filter's passband.

In the presence of hardware errors caused by voltage overscaling mentioned in Chapter 1.5, the conventional estimator is not optimal in terms of MSE anymore. Illustrated in Figure 2.1, a hardware error can cause a false peak,

which leads to a discrepancy between the estimated and the correct value of the input frequency.

Demonstrated ubiquitously in biological systems, the population activity of groups of neurons provides more accurate information than individual ones. This mechanism has particularly been observed in the control of eye and arm movements [13]. Based on these observations, a population-coding concept can be proposed and has two key features: 1) Correlated outputs of neighboring processing units are combined/fused to provide noise tolerance. 2) Low-precision and low-cost computations.

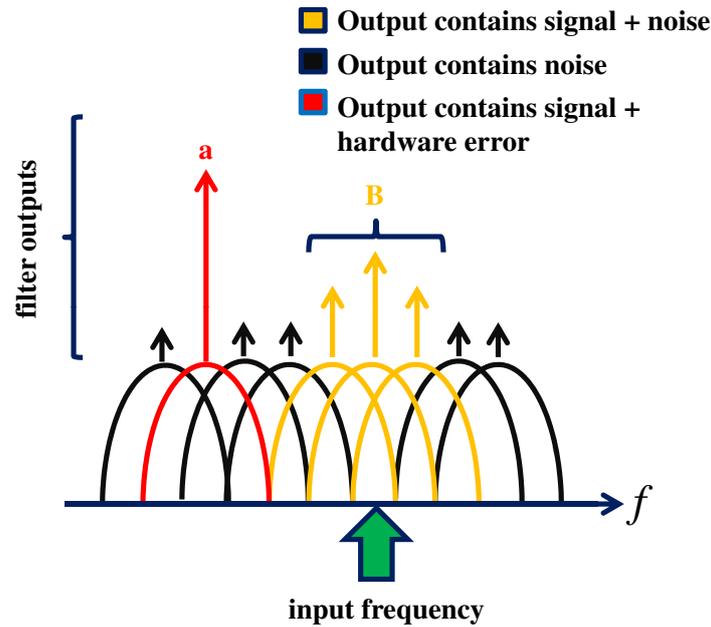


Figure 2.2: Overlapping filter with hardware redundancy. The arrow with ‘a’ indicates the filter output with hardware errors. The arrow with ‘B’ indicates a group of correlated filter outputs with the signal in the passband.

Inspired by this population-coding concept, a more error-tolerant solution is proposed for the problem imposed by the conventional non-overlapping filter approach. As shown in Figure 2.2, overlapping filters with wider bandwidth are used in combination with robust techniques (e.g., the median) applied to bundles of neighboring filters to achieve higher error tolerance. Instead of searching for the maximum of the filter outputs directly, the median value of the neighboring filter outputs is used to search the maximum. For instance, an isolated filter output with hardware errors (denoted by the

arrow with ‘a’) with two adjacent filters containing noise only will be ignored when the median estimator is applied to these three neighboring filter outputs.

2.2 The Low-Cost Sensor

Another benefit of using overlapping filters is that fewer taps are required to build FIR filters with wider bandwidth. In this work, coefficients for filters are obtained by using windowed linear-phase FIR digital filter design function (FIR1). MATLAB simulations show that 16 instead of 64 taps are needed to implement a filter with 1 kHz bandwidth instead of 0.25 kHz in order to keep the 20 dB reduction from mainlobe to sidelobe. This indicates significant hardware cost reduction and energy consumption savings, which can be related to the second feature of population-coding.

2.3 Signal Statistics in Time Domain under Voltage Overscaling

Besides frequency redundancy mentioned early in this chapter, signal statistics in the time domain under voltage overscaling can be also utilized to filter out the hardware timing errors caused by voltage overscaling [14]. As shown in the top panel of Figure 2.3, when the input frequency is not in the filter passband, with voltage overscaling the filter output in the time domain occasionally has errors, resulting in isolated abrupt signal amplitude change at certain time periods. In contrast, when the input frequency is in the filter passband, the filter output amplitude change over one time step could also be large. But because of the oscillating behavior of the filter output when the input frequency is in the filter passband as shown in the bottom panel of Figure 2.3, a filter output sample with similar amplitude might be found a few time steps earlier. Based on this property, a ANT filter can be developed to filter out hardware errors. The details are given later in this chapter.

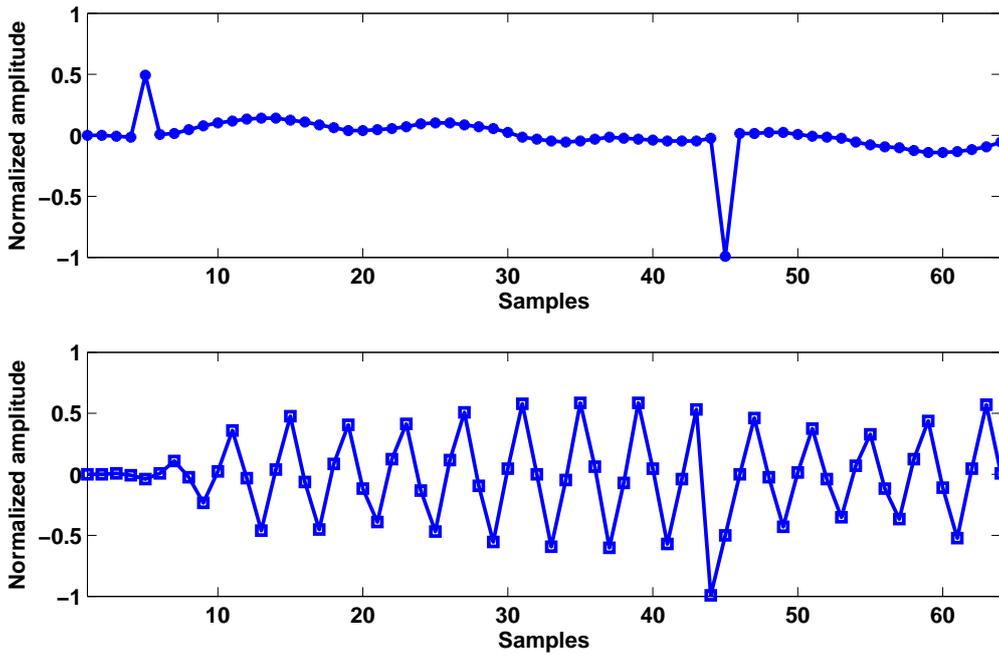


Figure 2.3: Filter outputs in the time domain with hardware errors.

2.4 The Nonoverlapping Estimator

As described earlier, the nonoverlapping estimator is only optimal in terms of MSE in the presence of additive white Gaussian noise and absence of hardware errors. Using the Hamming window design method, 64-tap FIR filters with a 0.25 kilohertz bandwidth are constructed. The filter with the maximum energy response indicates the value of the input signal frequency.

The block diagram of the conventional estimator is shown in Figure 2.4. The major components are the FIR filter banks (annotated by F_i), the energy estimator (annotated by E_i) and the frequency estimator (annotated by MAX). The individual FIR filter is implemented by the direct-form architecture with 8-bit inputs and 8-bit filter coefficients. The adders in the MAC have 22-bit outputs to avoid potential overflow. However, the final outputs to the next stage energy estimators only have 8-bit precision (the high 8 bits of the FIR filter output) as this is sufficient for the next stage computation. The energy estimator consists of a series of shift registers and records 64 samples in the time domain for a particular filter output. Then the maximum response within these samples is considered as the energy response.

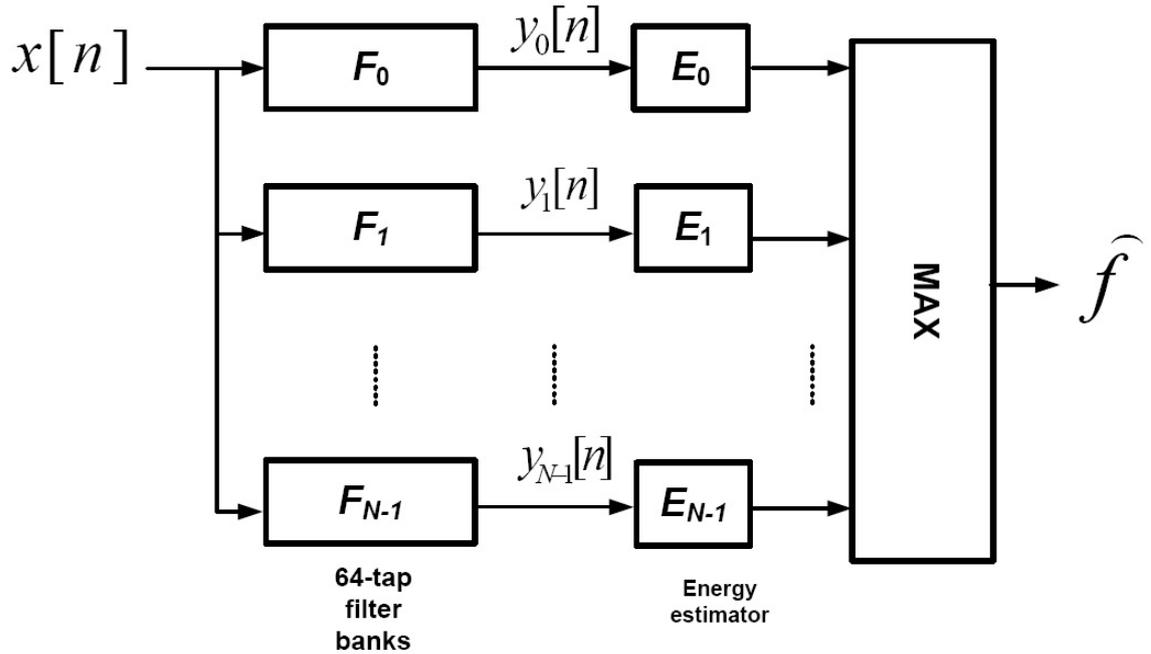


Figure 2.4: Conventional estimator overview.

Figure 2.5 shows the block diagram of the final frequency estimator based on the filter energy responses from the last stage. The elementary block in this diagram is a block with four sets of inputs such as the filter energy response and the index of one particular filter, which indicates the passband range of the filter. Its outputs are the max filter energy response of the four inputs and the corresponding filter index. Using this block, the estimator is implemented in a three-stage hierarchy [15]. The first stage has 16 basic blocks (annotated by M-A...M-R). Each takes four energy responses from the energy detector and passes the largest within the four inputs and the corresponding filter index to the next stage. The second stage utilizes four basic blocks (annotated by M1...M4) and generates 4 intermediate outputs in a similar fashion as in the first stage. The final stage use one basic block (annotated by M) and chooses the largest energy response and its filter index, from which the input frequency range can be derived (annotated by \hat{f} in the Figure 2.5).

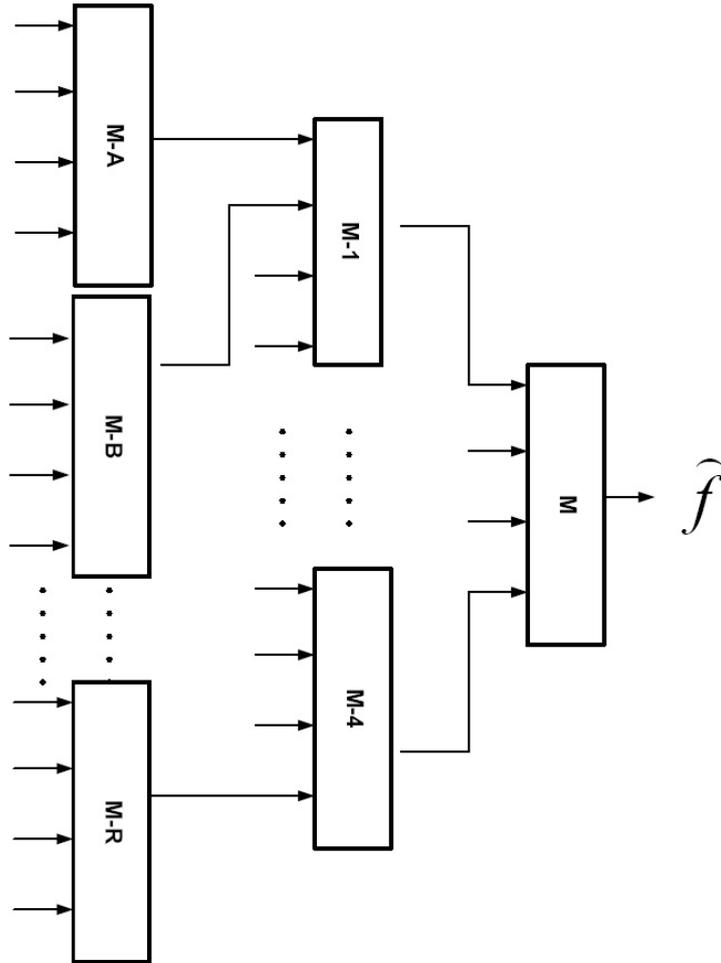


Figure 2.5: High level frequency estimator diagram.

2.5 The Bio-inspired Estimator

This estimator uses overlapping filters with wider bandwidth to achieve frequency redundancy. Using the Hamming window design method, 16-tap FIR filters with 1 kilohertz bandwidth are constructed while adjacent filters' bandwidth overlapping is 0.75 kilohertz.

The more detailed robust estimator is shown in Figure 2.6. It contains four major sub-blocks, including the prediction-based ANT filter bank (annotated by $P - ANT_i$), the energy estimator, the sliding median filter and the frequency estimator. The FIR filters F_i are similar to those in the conventional estimator. The difference is that 16-tap FIR filters instead of 64-tap filters are used. The input and the filter coefficients are both 8-bit. The adders in

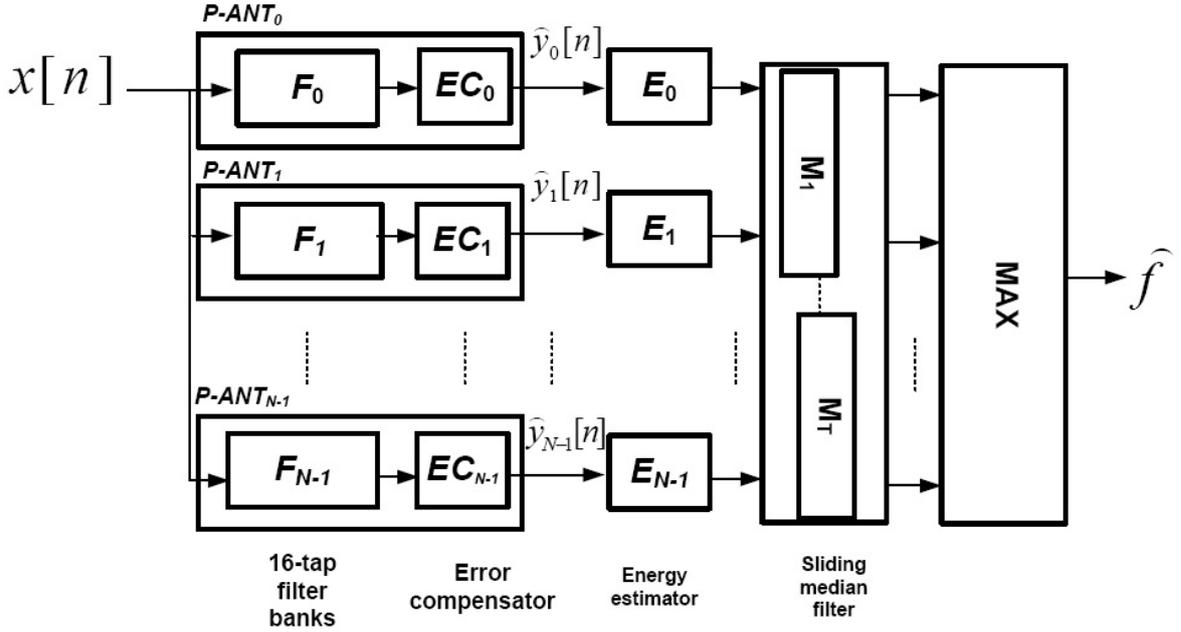


Figure 2.6: Robust estimator overview.

the MAC have 20-bit outputs instead of 22-bit as in the conventional estimator. The final outputs to the next stage energy estimators still only have 8-bit precision (the high 8 bits of the FIR filter output).

Figure 2.7 describes the structure of the ANT-based error compensator (EC), which takes the filter output from one of the upstream FIR filters and compensates for the hardware noise caused by voltage overscaling. A shift-register chain is used to store the filter output response of the last four time steps. If the signal magnitude change is larger than a quarter of the output dynamic range and there is no signal in the last four steps with magnitude similar to that of the current signal, then the current signal is replaced by the signal in the last time step. This means that the EC block employs a 1-step predictor to compensate for hardware errors [16]. The functionality of the EC block in Figure 2.7 is:

- Error detection: If $|y[n] - y[n-1]| > E_{th}$ and $|y[n] - y[n-2]| > E_{th}/2$ and $|y[n] - y[n-3]| > E_{th}/2$ and $|y[n] - y[n-4]| > E_{th}/2$, an error is declared. E_{th} is set as a quarter of the filter output dynamic range.
- Error correction: If an error is detected, $\hat{y}[n]=y[n-1]$; otherwise, $\hat{y}[n]=y[n]$.

The second sub-block in Figure 2.6 is the filter energy estimator and it is the same as in the conventional estimator described in Chapter 2.4.

The third sub-block is the sliding median filter, which takes the neighboring filter responses as inputs and passes the median energy response as the output to the next stage. For example, the i th median filter takes the group of energy responses E_{i-1}, E_i and E_{i+1} as inputs and chooses the median value as the energy response for the i th FIR filter. The next median filter takes the energy responses E_i, E_{i+1} and E_{i+2} as inputs and chooses the median value as the energy response for the next FIR filter.

The fourth sub-block is the frequency estimator, which takes the output of the sliding median block instead of using energy response from the energy response estimator directly, searches for the largest response and decides the input frequency range. The detailed architecture is the same as that used in the conventional estimator as in Figure 2.5.

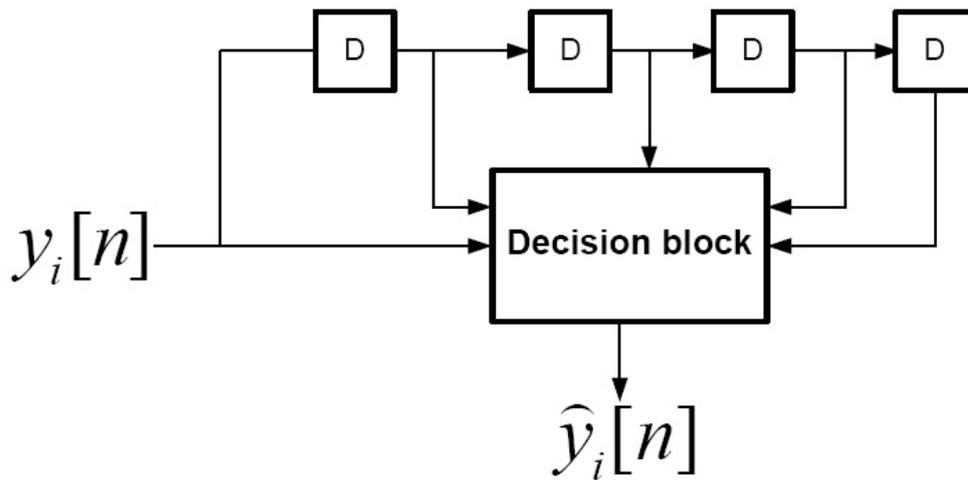


Figure 2.7: ANT-based EC block.

Chapter 3

Simulation Setup and Results Comparison

In this chapter, the performance in terms of RMSE is simulated using the RTL model of the algorithms. The energy consumption is based on the synthesis tool estimation (SYNOPSIS design compiler) and is normalized to total energy consumption of the conventional method under the nominal supply voltage.

3.1 Simulation Setup

In order to introduce timing errors to the FIR filter under different voltage overscalings, a structural Verilog model of the FIR filter is developed. As shown in Table 3.1, the sum and carry bit delay of the one-bit adder are simulated using SPICE in a 45 nm process under different voltage overscaling. Then an 8-bit ripple-carry adder (RCA) and an 8-by-8 bit signed multiplier are designed using the one-bit adder as the primary block. Finally, multiply/accumulate units (MAC) and FIR filters are constructed using the RCA and the multiplier.

The delay parameter in the adder's Verilog model can be modified according to different voltages. After the RTL model is constructed, the hardware error is introduced to each filter output by changing the delay parameters for specific $K_{vos} = V_{dd}/V_{dd-crit}$, where the $V_{dd-crit}$ is the critical supply voltage. Voltage overscaling is only applied to the combinational logic part of the FIR filter banks. The erroneous filter output is then fed into different estimation blocks under the nominal supply voltage to estimate the input frequency.

Table 3.1: Delay parameters for one bit adder driving a load of an identical adder under voltage overscaling

voltage(v)	Sum bit delay(ps)	Carry bit delay(ps)
1.20	45	41
1.15	48	43
1.10	51	46
1.05	55	48
1.00	61	52
0.95	68	57
0.90	77	63
0.85	90	70
0.80	100	80
0.75	121	95
0.70	151	115
0.65	230	150
0.60	381	221

3.2 Performance Comparison

3.2.1 Performance comparison

The performance of different estimators in terms of RMSE is summarized in Figure 3.1. RMSE is defined as:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (f_i - \hat{f}_i)^2} \quad (3.1)$$

where N is the number of trials, f_i and \hat{f}_i are input frequency and its estimated value, respectively, for a particular trial. The conventional method shows negligible but non-zero RMSE without voltage-overscaling because of the background additive white Gaussian noise. However, it exhibits no tolerance to hardware errors and its RMSE increases by two orders of magnitude under slight voltage overscaling.

At the nominal voltage and without hardware timing errors, the RMSE of the bio-inspired method is greater than the corresponding value of the conventional method, as the filter bandwidth is four times wider and there is more background Gaussian noise power leaking into the corresponding passband. Unlike the conventional method, the bio-inspired method shows

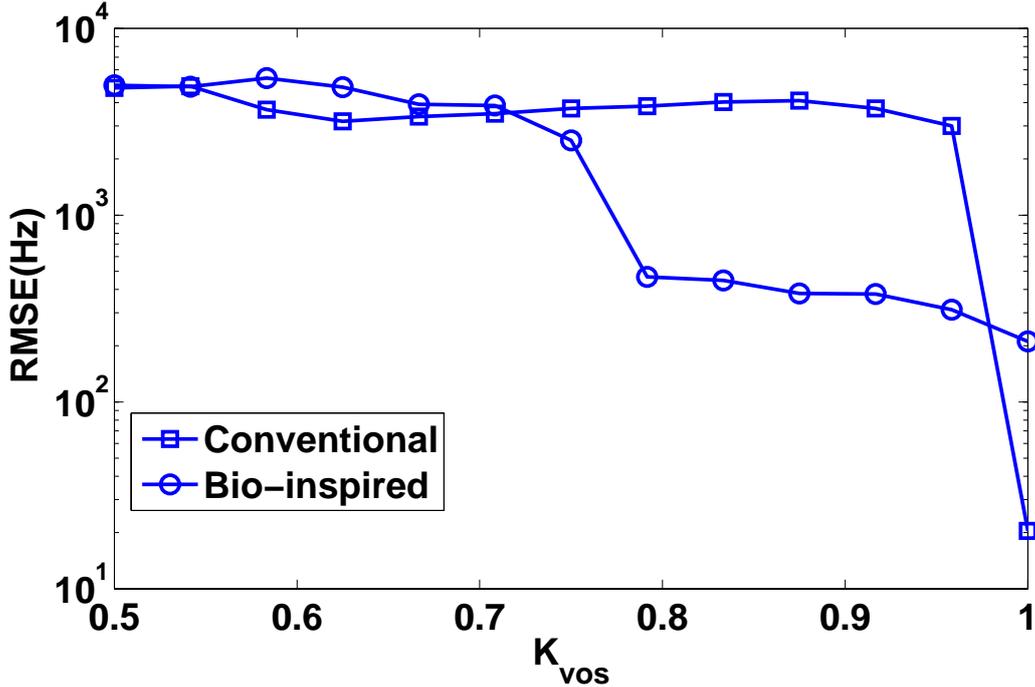


Figure 3.1: Performance comparison.

moderate tolerance against hardware errors introduced by voltage overscaling and the RMSE increases only slightly with further reduced supply voltage. When K_{vos} reaches 0.75, the performance of the bio-inspired method deteriorates greatly and the RMSE is at the same magnitude as in the conventional method.

3.2.2 Energy consumption comparison

It is necessary to assess the performance of different estimators in the context of energy consumption, which is estimated by the power and the operation time of different estimators. As no time redundancy is explored in this work, the operation time of different methods is the same. The energy consumption comparison would be equivalent to the power consumption comparison. The power consumption is estimated by the SYNOPSIS synthesis tool, ‘DC compiler’. In general, the system is divided into the FIR filter and the estimator. The power of each part is estimated individually and the total power is considered as the sum of the two parts. Under voltage overscaling, the

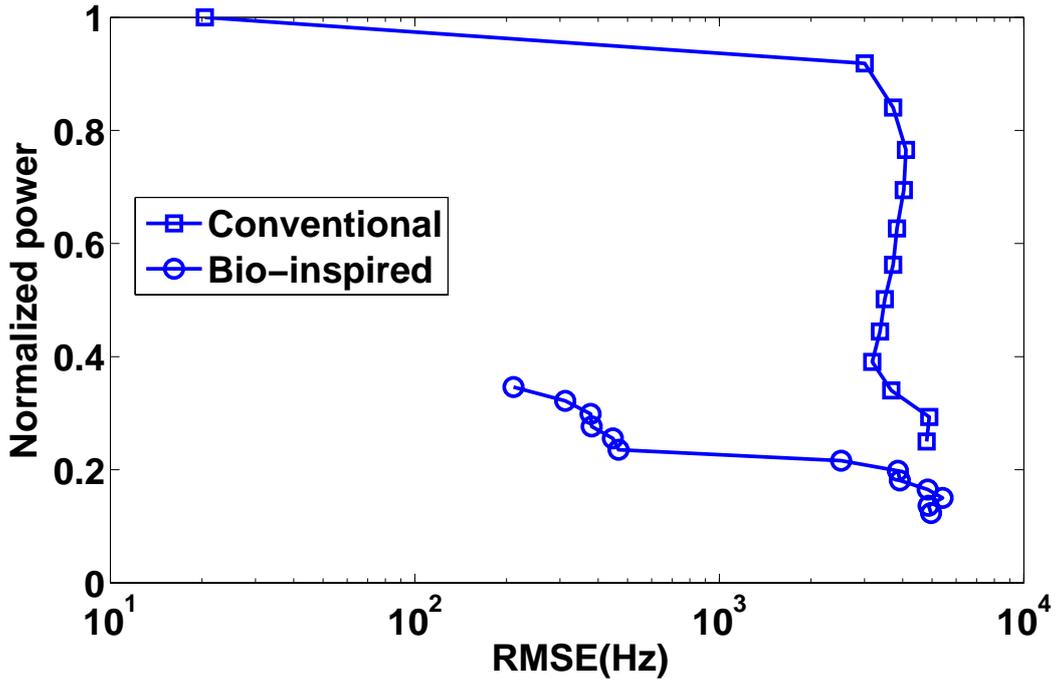


Figure 3.2: Power consumption at different K_{vos} .

power consumption of the filter banks is calculated by [14]

$$P = P_{V_{dd-crit}} K_{vos}^2 \quad (3.2)$$

Figure 3.2 summarizes the energy-consumption and performance trade-offs. According to (3.1), the conventional method's power consumption scales down approximately with the square of the supply voltage and exhibits no error tolerance as its RMSE increases by two orders of magnitude under slightly reduced supply voltage.

At the critical supply voltage, the total power of the bio-inspired method is 31% of the conventional method. The majority of the power saving can be attributed to the reduced hardware cost of the FIR filter banks. In the bio-inspired method, the FIR filter has 16 taps, which is only 25% of the tap numbers in the conventional method. The more complicated estimator introduces power consumption overhead. When voltage scaling is applied to the overlapping filter banks, more power saving can be achieved. However, this leads to deterioration of the system performance in terms of RMSE. At $K_{vos} = 0.79$, the bio-inspired method can reduce the power consumption by

up to 78% while keeping the RMSE under 500 Hz.

Chapter 4

Conclusion

4.1 Summary

Transistor feature size scaling along with subthreshold design provides significant energy consumption benefits for many sensor network applications such as security and environmental sensing, which often take place in a highly dynamic and variable environment. However, implementations using smaller transistor size or operating under the subthreshold voltage often introduce variations that can cause timing errors.

In this thesis, a novel algorithm for non-coherent frequency estimation of a sinusoid in noise inspired by biological signal processing systems is developed. This method is shown to be energy-efficient and error-tolerant for the canonical problem addressed.

The bio-inspired method can reduce the RMSE to 500 Hz compared to the conventional method results of 5000 Hz RMSE. Moreover, the bio-inspired method's energy consumption is only 22% of the conventional method.

In this work, hardware redundancy is mainly introduced by using low-cost 'sensors' which have overlapping passband and lower energy consumption compared to the conventional design.

This method is mostly effective under weak or moderate voltage overscaling. Under this condition errors happen occasionally.

4.2 Future Work

This work mainly has explored the frequency redundancy applied to this canonical single-tone audio frequency problem. Further research can be expanded in several aspects. First, we could explore redundancy in time, which utilizes more samples in time for the same input frequency. This introduces extra energy consumption while it could improve system performance

in terms of RMSE. Energy/performance trade-offs are further explored by altering the number of iterations.

In this work, hardware errors are introduced by deliberately overscaling the supply voltage of the digital filter part of the system. Similarly, hardware errors could also be introduced by other variations such as process, temperature and aging. Other hardware error models such as Stuck-At fault can also be explored. For the fusion part of the system, alternative architectures for ANT filters can also be explored and the length of the ANT filters can be optimized for different input frequencies.

Further research can expand the bio-inspired method with features such as low-cost rough detectors/estimators with redundancy into other detection and estimation problems mentioned in Chapter 1.4 to design energy-efficient and error-tolerant algorithms for certain applications. To achieve this, the key is to design low-cost and energy-efficient detectors/estimators compared to the conventional design. Further, hardware redundancy can be introduced given the low-cost detectors/estimators are available.

References

- [1] J.M. Rabaey, D. Burke, K. Lutz, and J. Wawrzynek, “Workloads of the future,” *IEEE Design and Test of Computers*, July-August 2008, vol. 25, pp. 358–365.
- [2] J.M. Rabaey, A. Chandrakasan, and D. Nikolic, *Digital Integrated Circuits, A Design Perspective*, Englewood Cliffs, NJ, 2002, Prentice-Hall.
- [3] C. Hu, “IC reliability simulation,” *IEEE Journal of Solid-State Circuits*, March 1992, vol. 27, pp. 241–246.
- [4] K.-K. Kim and Y.-B. Kim, “A novel adaptive design methodology for minimum leakage power considering PVT variations on nanoscale VLSI systems,” *IEEE Transactions on Very Large Scale Integration Systems*, April 2009, vol. 17, pp. 517–527.
- [5] J.M. Rabaey, *Low Power Design Essentials, Integrated Circuits and Systems*, New York, NY, 2009, Springer.
- [6] N. Verma, J. Kwong, and A. Chandrakasan, “Nanometer MOSFET variation in minimum energy subthreshold circuits,” *IEEE Transactions on Electron Devices*, January 2008, vol. 55, pp. 163–174.
- [7] B. W. Johnson, *Design and Analysis of Fault Tolerant Digital Systems*, Boston, MA, 1988, Addison-Wesley Longman Publishing Co.
- [8] J.S. Plank, K. Li, and M.A. Puening, “Diskless checkpoint,” *IEEE Transactions on Parallel and Distributed Systems*, October 1998, pp. 972–986.
- [9] D. Ernst, N. Kim, S. Das, S. Pant, R. Rao, T. Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner, and T. Mudge, “Razor: A low-power pipeline based on circuit-level timing speculation,” in *Proceedings of the 36th International Symposium on Microarchitecture*, November 2003, pp. 7–18.
- [10] S. Narayanan, G. Varatkar, D. L. Jones, and N.R. Shanbhag, “Computation as estimation: Estimation-theoretic IC design improves robustness

- and reduces power consumption,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing (ICASSP)*, April 2008, pp. 1421–1424.
- [11] G. Varatkar, S. Narayanan, N.R. Shanbhag, and D. L. Jones, “Variation-tolerant, low-power pn-code acquisition using stochastic sensor NOC,” *IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2008, pp. 380–383.
 - [12] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*, Englewood Cliffs, NJ, 1998, Prentice-Hall.
 - [13] W. H. Rohrer, C. Lee, and D. L. Sparks, “Population coding of saccadic eye movements by neurons in the superior colliculus,” *Nature*, 1988, vol. 332, pp. 357–360.
 - [14] R. Hegde and N. R. Shanbhag, “A voltage overscaled low-power digital filter IC,” *IEEE Journal of Solid-State Circuits*, February 2004, vol. 39, pp. 388–391.
 - [15] E.P. Kim, D.J. Baker, S. Narayanan, D.L. Jones, and N.R. Shanbhag, “Low power and error resilient pn code acquisition filter via statistical error compensation,” *IEEE Custom Integrated Circuits Conference (CICC)*, October 2011, pp. 1–4.
 - [16] R. Hegde and N. R. Shanbhag, “Soft digital signal processing,” *IEEE Transactions on Very Large Scale Integration Systems*, December 2001, vol. 9, pp. 813–823.