Custom Silicon for Low-Cost Information Dissemination among Illiterate People Groups

by

ZhiYoong Foo

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering)
in The University of Michigan
2013

Doctoral Committee:

Professor David Blaauw, Chair
Professor Peter Chen
Assistant Professor James Cutler
Assistant Professor Prabal Dutta

To my loving family.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER 1

Introduction

1.1 Thesis Scope

In this thesis, we present the work done at the University of Michigan where the improvement of the quality of life of the poorest people in the world is attempted through Very Large Scale Integrated Chips (VLSI). The premise of the work was to see if we could build an ICT device that can be used to disseminate information to and among the poorest people in the world, identified to be subsistent farmers living in rural agrarian societies in developing regions. We first show how the benefits of timely, actionable information can improve the quality of lives. We then go on to show how current NGO solutions that attempt this have shortcomings and how current ICT solutions, while successful in some aspects, fall short in tackling all the unique problems in this setting: cost, power, connectivity, usability, robustness and illiteracy. We argue that even when an ICT solution is designed specifically to overcome the unique problems presented, that it falls short due to the high price and the low purchasing power of the target population.

We show that by designing custom silicon targeted to this application space, we can overcome the challenges presented. Literacy in Technology (LIT), developed at University of Michigan, tackles all the challenges through design decisions such as having a high level of integration to reduce the initial cost to a point that can be purchased by the target population. Other side effects of this design decision are a lowered Printed Circuit Board (PCB) footprint size due to the lowered count of passives and reducing the number of on-

board chips to only 3. Furthermore we show how by reducing the costly DRAM or NOR Flash traditionally used for code storage in microcontrollers and leaving both code and data in an on-board NAND Flash further reduces the cost of the ICT solution. Other decisions such as removal of wires and open ports and using wireless delivery of content through radio updates or near field inductive communication improves the robustness of the solution. We also ensured that we designed LIT to consume as little energy as possible thus lowering the recurring costs of the end users. We also store information in an audio based format so that literacy is not a requirement for using the device. Other design decisions such as reusing traces that will be used during the PCB manufacturing for touch sensors and inductive coil increases the robustness of the design while reducing the cost as well. Through such design decisions, we show that LIT can overcome the unique problems presented.

1.2 Information

Information, knowledge that can be derived from study, experience or instruction, is widely available to most people in developed regions. From local libraries, to radios, televisions, and the internet, most people in developed regions can access information easily. Powered with information, people can be educated about basic rights, public services available to them, better healthcare, education, and work opportunities, all which can drastically improve their quality of life. Access to information is a resource that most people take for granted. However, a large number of the worldwide population suffers from difficulty in accessing information either due to reasons such as low purchasing power and illiteracy. Their low purchasing power means that they have limited capabilities in choices of how they can get access to information. Furthermore, literacy allows people to be independent in their means of access to and absorbing new information and passing of knowledge between groups through reading and writing in print. Once only restricted to a small elite group of people, literacy has been shown to have a large impact on people's livelihoods.

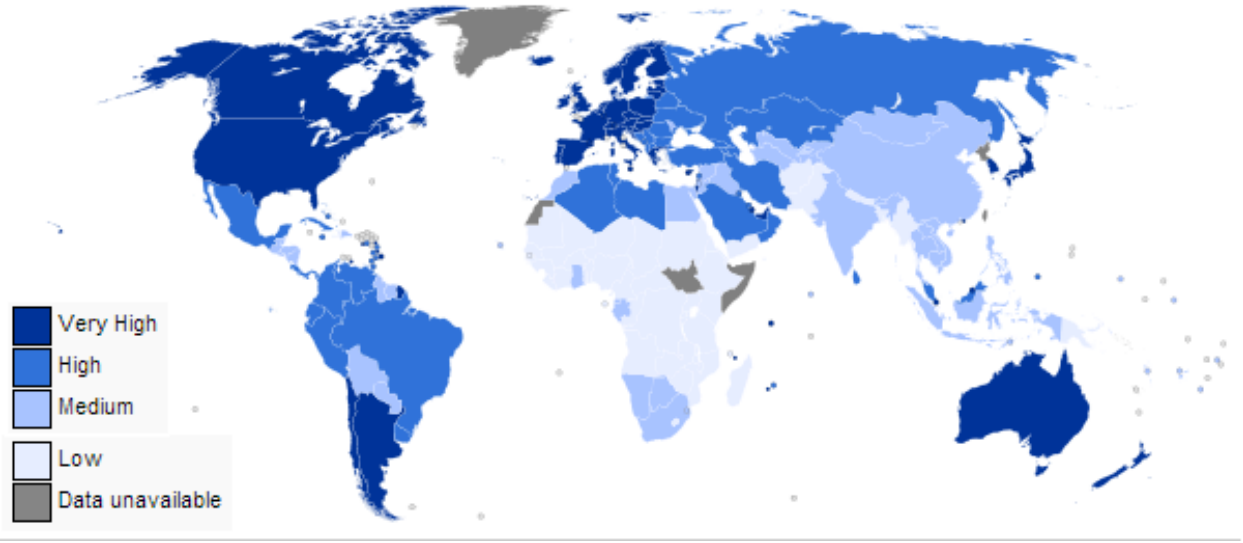

Figure 1.1: 2013 Human Development Index World Map

1.3 Quality of Life

The Human Development Index (HDI) is a composite statistic of life expectancy, education, and income used to rank countries and the quality of life of its citizens [2]. We can see from (Figure 1.1) that developing countries generally have low HDI scores. These developing countries have a high percentage of the population living in rural settlements where access to information is difficult. These populations rely on subsistence farming, a practice where farmers grow enough food for themselves and their families. Furthermore these population groups have a high number of illiterates making improving their quality of life through information difficult.

1.4 Impacts of Information, Literacy, and Education

Recent studies have illustrated the benefits of timely, actionable information on rural agrarian populations in developing regions. In the Avaaj Otalo project [3], a mobile phone-based, interactive voice forum enabled the "emergence of norms, persistent moderation, and a desire for...interaction with...authorities and open discussion with peers" in India. Users of the system found value in listening to questions and answers with 77% of interviewees citing this as the main reason they liked the forum. The Digital Green project [4], "seeks to

disseminate targeted agricultural information to small and marginal farmers in India using digital video." A four month trial of the system showed a six to seven times increase in adoption of certain agricultural practices. The Talking Book project [5] deployed a custom-built audio computer in rural Ghana to study "the impact of giving on-demand access to guidance created by local experts" and found that in a village with 90% illiteracy and no electricity, Talking Book users significantly increased crop production over farmers who did not use the devices.

Literacy and education have also been shown to correlate with issues ranging from poverty, to incarceration, to high health care costs [6] in the United States. For persons in developing regions, the same issues reign, while the consequences of illiteracy have also been linked to child mortality rates [7]. Children, whose mothers can read or write have a higher chance of surviving since they are more knowledgeable about health and nutritional practices and live in better situations. Child mortality has also been shown to reduce by an additional 8% for every year that a mother is educated [8]. These studies show how access to information can directly impact quality of life and is a motivator to finding means to supply information to those that need it the most: subsistence farmers living in rural agrarian populations that have a high number of illiterates.

1.5 Thesis Overview

In the following chapters, we dive into the details of the work done for the thesis. Each chapter will highlight the motivation behind the work, talk about current solutions and present how LIT overcomes the challenges presented to achieve its target goals.

Chapter 2 explores the unique challenges of dissemination information amongst illiterate persons in developing regions. We show how current NGO solutions have limitations that can be overcome through ICT solutions. We then showcase some examples of current ICT solutions that attempt to tackle the problem and point out their shortcomings. By highlighting their shortcomings, we see that there is a real need to design a custom solution specifically tailored to fit the needs of the targeted end users. One such custom solution, the Talking Book, showed encouraging results from its pilot study, but it is still at a price point

that is not viable for the end users that subsist on $1 - $2 a day. We show that in order to overcome all the complexities, custom silicon will yield a result that can cut the Gordian Knot that ties cost, power, connectivity, usability, robustness and illiteracy.

Chapter 3 introduces LIT and provides a quick overview of its unique capabilities. LIT shows how judicious design of silicon can improve upon previous ICT designs and can overcome the challenges — cost, power, connectivity, and robustness posited by the task of disseminating information among illiterate people. With its high-level of integration, novel circuit level, system level, and application specific level solutions and design decisions, they all work in unison towards the achieving the primary goal of reducing its initial cost and recurring cost to a point that can be purchased and used by the end users.

Chapter 4 highlights LIT's entirely On-Chip Integrated Power Management System. Resulting in a low off-chip component count that allows the final ICT solution to be priced at a point viable to the end users, we show how commercially available solutions do not have the level of integration that LIT achieves. This high level of integration differentiates LIT from current ICT solutions that are too expensive thanks to having a high number of off-chip components. Combined with LIT's Power Operational Modes, Active, Sleep, and Deep Sleep which results in LIT's 2 year lifetime. LIT's Board Level Configurable Pads allows LIT to be flexible in choice of either a 1.8V or 3.2V off-chip NAND Flash chip allowing LIT to be sensitive to market price volatility of the NAND Flash chips. This also results in a flexible low initial cost solution that other ICT solutions do not manage to achieve.

Given LIT's low-cost goal, Chapter 5 shows how LIT has to provide power to a whole system entirely from a single chip while maintaining a low off-chip component count which will result in a low-cost in order to power its off-chip components. We argue that inductors commonly used in Boost and Buck converters are not feasible, namely due to its cost and on-chip design complexity. Switch Capacitor Networks are the best solution due to the low cost of capacitors. However, they suffer from fixed conversion ratios and since LIT is designed to operate on a wide battery input voltage, we achieve generating the 3.2V through the design of a novel Hybrid Switch Capacitor Network that combines a Step Down Switch Capacitor Network that provides fractional values that gets fed into a Step Up Switch Capacitor Network. This allows us to achieve the 10% range of 3.2V in order to power the

off-chip components while still maintaining LIT's design goal of a low-cost ICT device.

Chapter 6 talks about the investigative work behind currently available commercial microcontrollers and other possible memory solutions that led us to build LIT's memory hierarchy. We found that current microcontrollers are either too expensive as a one minimal solution system or severely lacks the capabilities that we require if we wanted to meet our target cost goals. Other possible memory solutions were either too expensive (MicroSD, NOR Flash, DRAM) or unfeasible to be implemented on chip (Embedded Flash). By using a 128kB 4-way True LRU cache backed by a large off-chip NAND Flash chip in conjunction with a low power, low cost small microcontroller, this allowed us to meet our target cost goal. However, this decision led to some complexities and obstacles that we managed to overcome. We also show how LIT's memory reduces its energy consumption through granular power gating that both allows us to maintain flexibility in the amount of code left awake for the event validity checker that reduces false wake ups.

Chapter 7 shows how LIT's novel robust Power-on-Reset / Brown-out-Detector with lock-off feature with dual comparators and hysteresis manages to overcome the oscillations due to the Carbon-Zinc batteries used in these developing regions that have a strong self-healing property. By not allowing reset to deassert if LIT had already been in reset when the Carbon-Zinc battery voltage drops below 1.7V, we manage to not corrupt LIT's NAND Flash's contents. Furthermore by having the Power-on-Reset / Brown-out-Detector with lock-off only allowed to be reset past 1.2V, we ensure that with such a large margin of hysteresis, the self-healing properties of the Carbon-Zinc batteries will not cause LIT to oscillate between power-on and off.

Chapter 8 showcases how LIT is sensitive to robustness and cost through the overloading of PCB traces used as both a human input interface device and a component in data transmission between devices. By reusing the traces already available on PCBs during the PCB manufacturing process, LIT manages to keep cost down by not requiring additional components to implement the human input interface and data transmission off-chip. Furthermore, the traces improve the robustness of the system by removing the moving parts such as mechanical buttons found in the previous ICT solutions which are a failure point. It also removes the open ports, dongles and wires required to transfer data from device to

device further improving upon previous ICT designs. Through the careful reusing of components already available during the manufacturing process, we show how LIT can achieve its goals of a low initial cost ICT solution.

Chapter 9 presents LIT's results attained from silicon testing with Core Frequency vs. Voltage, silicon measure core GOPS/WATT, Shmoo plot, design parameters and their associated measured results, the expected Bill of Materials table broken down into the total chip cost, board cost and system cost and its associated components along the way.

CHAPTER 2

A Case for Custom Silicon in Enabling Low-Cost Information Technology for Developing Regions

2.1 Motivation

In Chapter 1, we highlighted the importance of information and the benefits of timely, actionable information through studies that show how access to information and literacy can improve not just economical status, but also healthcare. Given the importance and the impact of access to information on the livelihoods of the nearly 800 million who are illiterate and have difficulty accessing information, solutions must be identified to overcome it. Encouraged by results illustrating the benefits of timely, actionable information on rural agrarian populations in developing regions, the following sections explore the question of how to scale, and make economically viable and sustain, such technologies for the nearly 1.5 billion people who live without electricity, the nearly 800 million who are illiterate, the great majority of whom do not own mobile phones and live in rural areas with limited or no data connectivity [9][10].

First, we will explore the practical challenges for information dissemination in rural settings in Section 2.2. We show that the current method for transferring knowledge to persons in rural settings is inadequate and that a new methodology is required. We argue that cost is the most important factor that needs to be addressed in order for the new methodology to be adopted in a widespread, sustained deployment and be successful. Furthermore, we

present the practical challenges faced by designers that are interested in tackling this problem from an Information and Communications Technology (ICT) perspective, specifically: cost, power, connectivity, usability, robustness and illiteracy.

In Section 2.3, we outline why existing devices, most notably mobile phones, have been unsuccessful in meeting this need. In particular, we show how smart phones fall significantly beyond the required cost bracket. Even inexpensive basic phones do not quite address the needs of this population. Basic phones, while approaching the required cost point, lack the data storage, offer no simple data transfer methods, and pose logistical issues in that they require "reflashing" (an operation highly-specific to each phone model). In addition, mobile phones have high power draw and poor speaker efficiency, which makes their ongoing cost of ownership significantly higher than the devices currently owned by these populations. Other currently available Information and Communications Technology (ICT) solutions all have significant shortcomings as well. Given the shortcomings of currently available solutions, we argue that developing a device designed to tackle the unique challenges presented is needed in this application space.

In Section 2.4, we introduce Literacy Bridge's Talking Book, a device that has been developed using commercial off-the-shelf (COTS) components to tackle the unique challenges presented in this application space. The Talking Book was specifically designed for information dissemination amongst illiterate persons and we cite its successes in its pilot study. However, the Talking Book is not without its shortcomings as well, and we argue that its disadvantages can be overcome with custom silicon instead of a COTS design.

Given the need for a custom device, in section 2.5, we present the Rural Audio Computer, a device tailored to the unique Information and Communications Technology (ICT) needs of rural developing regions and the system architecture that supports its operation. We then argue that custom silicon offers a way to cut the Gordian Knot that ties cost, power, connectivity, usability, robustness, and illiteracy. Through: (i) component integration, (ii) reduced power draw, (iii) application-specific design optimizations, and (iv) electronic-mechanical co-design, a custom silicon design for Rural Audio Computers offers dramatic cost reductions that can make information technology widely accessible.

2.2 Practical Challenges for Information Dissemination in Rural Settings

System designers using ICT targeted at developing regions face many different challenges than those that apply to more affluent markets. In this section, we outline some of these challenges and discuss contemporary mechanisms for information access. For the world's poorest 2.6 billion people that live on $2 per day or less and make up approximately 40% of the earth's population that lack literacy, disseminating knowledge to them in a viable manner has many unique and practical challenges that need to be overcome.

Currently, the most efficient way to teach practices for reducing child mortality rates is often by driving a pickup truck over rough unpaved roads to a remote village that hasn't been visited recently. Upon arrival in a targeted village, the local health expert might gather anyone who is available and overwhelm them with numerous messages about disease prevention (Figure 2.1). Unfortunately, those who are able to attend will have difficulty recalling any specific guidance that they did not immediately apply — they are not able to replay the guidance on demand nor can they take notes due to lack of literacy. Furthermore, those who did not attend will suffer from the "telephone game" effect of having information spread through word-of-mouth from those who attended, diluting the guidance given. There needs to be a solution where they can retain the information being passed to them and access it, on demand, in a format that does not rely on literacy. This current method for transferring knowledge is clearly inadequate and we argue that ICTs can yield a solution.

Finding an ICT solution for this rather unique scenario has multiple challenges — cost, power, connectivity, usability, robustness, and illiteracy. We argue that cost is the least flexible challenge facing timely access to actionable information for illiterate persons. If cost were not a factor, one could conceivably provide audio information to illiterate persons using audio recordings on a smart mobile phone and simply recharge batteries using a solar panel or generator. One could use GPRS or EDGE, where available, for content downloading or use an integrated or external FM receiver to download content where GPRS or EDGE was not available, share content using WiFi in ad hoc mode, and offer graphical or voice-activated user interaction. The technology for this is readily available, proven, and can be

Figure 2.1: Villagers gathered around a tree attending tutorial from a visiting NGO.

easily implemented if cost were not an issue. For designers of ICT products aimed at this population, cost is key. Designers must consider both the cash-flow challenges associated with initial purchase and the overall return on investment, which factors in all ongoing costs. The total cost of ownership (TCO) that designers must be aware of include both the upfront cost, and recurring cost.

That upfront cost is the dominant factor in widespread, sustained deployment is illustrated by two first-hand observations from a recent trip to remote villages in Ghana by members of our research team. First, these subsistence farmers (who grow enough food for their families with little or no surplus for cash income) tend to purchase a low-quality ax that costs $10, but usually breaks quickly and needs to be replaced, rather than spend money on a high-quality ax that costs $20 but will last a life time. This indicates that they tend to not be able to save the capital up to invest in something that will yield better returns. Second, the majority of electronics owned by these populations, again based on our investigations, typically cost less than $10 and are almost never more than $20, establishing the viable price points for typical consumers: economic viability requires a target price of approximately USD $10.

The recurring cost, overhead of accessing information from the device, must also be considered. For example, consuming large amounts of energy to retrieve the information would make the device expensive to use by requiring end users to have to repeatedly recharge or buy new batteries. While lowering the energy consumption of the device is an overhead that designers must be aware of, access to content must also be considered. For instance, a basic mobile handset may cost between USD $20-$40, but listening to new content costs an additional $0.10-$0.30 per minute throughout most of Sub-Saharan Africa [11]. From our investigations, we also observed that the end users often prefer to listen to a recording several times. Given that, we can conclude that this recurring cost could quickly dominate the TCO making using the device not unfeasible.

While we argue that both upfront and recurring costs are the least flexible challenges, designers need to be aware of the other challenges as well — power, robustness, connectivity, usability, and illiteracy. Power is a concern since worldwide, 1.5 billion people have no access to electricity and, what's worse, is that the situation is not expected to improve much in

Figure 2.2: The world at night: Showing electricity penetration.

the next 20 years: by 2030, the International Energy Agency predicts the number to drop from 1.5 to 1.3 billion people Figure 2.2). Only 15% of rural households in Sub-Saharan Africa have access to grid electricity [12]. Therefore, any ICT product designed for the poorest people in the world cannot depend on grid electricity. Similarly, designing a device to include a built-in renewable power source will reduce the TCO in the long-run through power cost reduction, but the additional cash required to acquire the device may put it out of reach for this audience. For example, Global Recordings Network's Saber [13] and One Laptop Per Child [14] are such devices that allow the users to recharge them but puts a much higher burden on the upfront cost. Designers also need to ensure that the ICT solution is robust enough to be handled by the end users who may not be as careful with the device as typical users that are more accustomed to technology. Moving parts that are used to recharge the internal batteries will be prone to mechanical failure and will have to be designed to be more robust and thus burden cost even more.

Connectivity and usability are also challenges that need to be considered from this unique perspective. The human-computer interface needs of the end users are significantly different from literate users [15]. ICT designers should also consider efficient ways of delivering information to the end users through the device. Studies of constant pitch speed shifting of audio

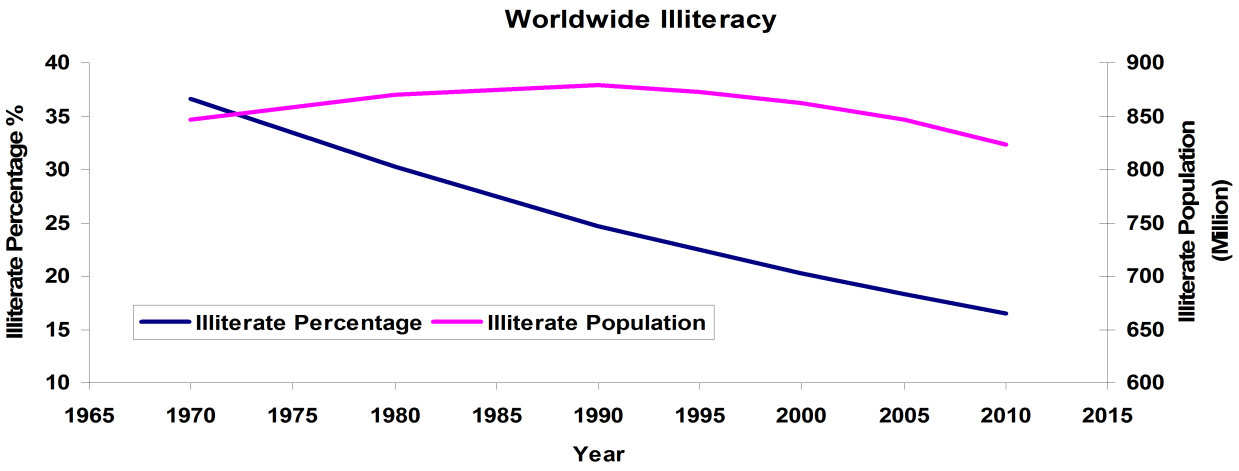

Figure 2.3: Worldwide Literacy Rates: Percentage of illiterate people is decreasing, but total illiterate population is remaining constant

have shown that listening to materials sped up at least twice is more effective than listening to it at normal speed [16]. Furthermore, constant pitch speed shifting allows users to slow down or, scan and skip sections in a manner similar to reading printed text. This helps with understanding the material since the end users can pace it an understandable rate [17].

According to the UNESCO's Institute for Statistics, the total percentage of literate adults (adults defined as 15 years and older) in the world is 84% with 88.6% of the males and 80% of the females literate [1]. While the total percentage of illiterate persons have decreased over the years, the total number of illiterate persons has remained roughly constant (in part due to population growth) (Figure 2.3). 775 million adults, of which 64% are women, are unable to read or write. Literacy rates are the lowest in sub-Saharan Africa and in South and West Asia (Figure 2.4).

Across these regions, self-reported adult literacy rates are approximately 60% (these numbers are typically self-reported and have been shown to significantly underestimate the problem when follow up testing occurs). In many villages, the literacy rate may be as low as five to ten percent. Total youth (youth defined as ages 15 to 24 years) illiteracy is at 122 million. The nearly 800 million adults in the world that are illiterate have difficulty gathering information the way that the majority of the world's population can. This inability to gather information affects their ability to educate themselves and presents a unique problem to information access from an ICT solution perspective.

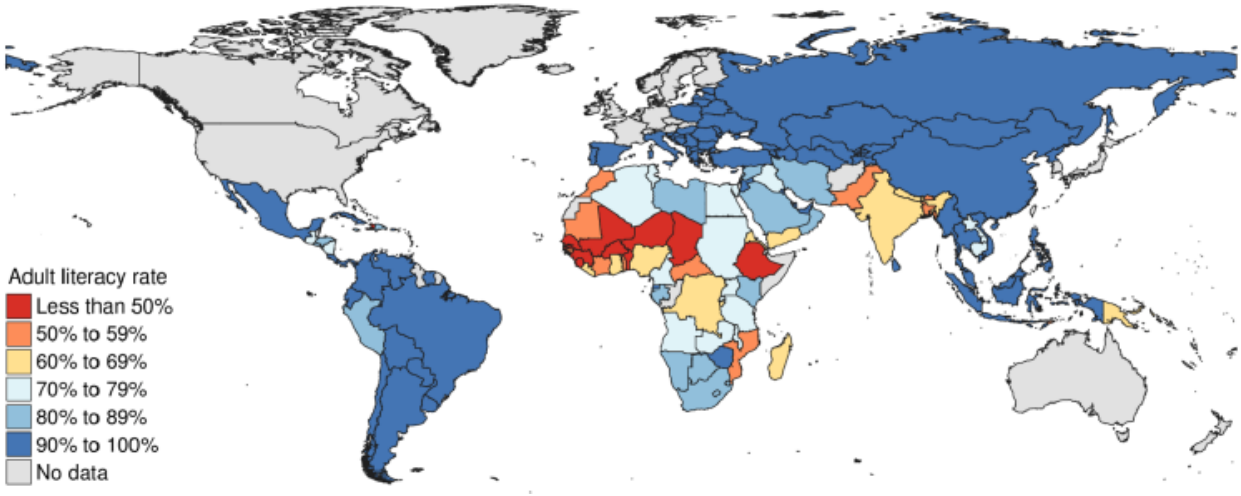

Figure 2.4: Global Map of Literacy Rates for Adults 2010 [1]

2.3 Contemporary Dissemination Vehicles and their Limitations

In this section, we showcase a few examples of current ICT solutions that attempt to meet the needs described in Section 2.2. From historical ICT solutions such as the radio, to state-of-the-art solutions such as mobile phones, we highlight how certain aspects of the current solutions work to solve the challenges. We also discuss their shortcomings that fail to make their widespread, sustained deployment successful.

Radio — Radios are widely used and available throughout rural developing regions and are at both an initial and recurring cost point that end users can afford. From our research team's observations in Ghana, a radio typically costs $5. However, its limitations are that it does not allow end users to find the information they need, nor does it allow users to listen to a message several times at times convenient to them since they need to physically be there while the information is being transmitted. This is important since end users typically listen to a message multiple times based on our observations (Figure 2.5).

Cassette Players — Cassette Players are also available in rural developing regions and are at an initial cost point that end users can afford. However, with their electro-mechanical action, they are not energy efficient and add significant costs for content reproduction, distribution, and sharing with more than one device. Their electro-mechanical action also lowers

Figure 2.5: Typical radio available in developing regions.

Figure 2.6: Typical cassette player available in developing regions

their robustness as it will be prone to mechanical failure (Figure 2.6).

Cellular Phones — Cellular phones have become a necessity for nearly everyone in developed regions (Figure 2.7). The ability to have information at the touch of your fingers almost anywhere you go makes cellular phones a very tempting ICT solution. A cellular phone can download, store, and recall an audio recording which allows users to replay content as necessary, but there are many advantages and disadvantages to cellular phones which we will discuss. First, despite the cost of cellular phones dropping, cellular phones are still priced at roughly $20, based on our research team's first hand observations in Ghana. This requires an initial investment that is just not possible for most of the poorest 40% of the world. Second, while only requiring power as the cost of access to the content stored on the phone, cellular phones need to be recharged, and some cellular phone owners will walk two or three hours to recharge their phone in a nearby town since their own town lacks the

running electricity as we have shown in 2.2. This scenario of crossing to other towns only work best when the phone is used for infrequent and brief calls. In fact, phone usage may be infrequent by necessity when network reception is too weak or not available in one's village (as applies to millions of mobile subscribers).

Connectivity is one of the challenges that a cellular phone, at first glance, might prove to be an advantageous ICT solution. Basic phones in Africa have a SMS cost between $0.05 to $0.21 [18] which is comparable to a 1 minute voice call [11]. SMS feedback is short and small compared to the data volume necessary to receive audio data and thus, SMS can provide a high-cost, low-bandwidth feedback channel. Another possibility of retrieving specific content on the phone is through the GSM data channel. The General Packet Radio Service (GPRS) or Enhanced Data rates for Global Evolution (EDGE) allow direct download of digital audio data over the Internet to a phone. However, these services are not available everywhere, and require an additional data plan. Access to the data modem of the phone is another issue that the solution would have to overcome: not all phones expose this capability to external devices, and there is no standard modem port on cellular phones. A further possibility for direct content dissemination is through the GSM voice channel. Recent work provided methods for transmitting digital data over the GSM voice channel [19][14] show that this is possible. While a cellular phone's connectivity basic phones and networks in the most impoverished areas often include SMS text capabilities, and sometimes even digital data services, there is a significant barrier for most potential users: illiteracy. Any ICT solution depending on text is not an option for the people who have the greatest need.

In each of these scenarios, a successful example of using mobile phones for communication does not necessarily translate into the ability to use the same tool for ongoing learning and on-demand reference. Some mobile phone-based ICT solutions have adopted interactive voice response (IVR) systems. The main challenge with such a system is usability, particularly on the input side (user selects an option). Those systems applying voice recognition are limited to the largest languages, which again excludes most of the potential beneficiaries of such a system. Those systems using keyed input face a usability challenge, particularly for people without numeracy skills or without exposure to any ICT beyond a radio. On the output side (prompts and status messages), graphical icons increase device cost and exclude

Figure 2.7: Typical basic cellular phone available in developing regions.

those with a visual disability; visual disability is twice as common in developing countries as developed countries — and yet Braille is much less available and not supported by cellular phones. Mid-range phones with FM radios and recording capability are still too expensive for one billion people living in extreme poverty; but more importantly, they are not designed as learning devices for people raised in oral cultures with little formal education.

Global Recordings Network's Saber — The Saber [13] does not require literacy to be operable and has a hand crank that can be used to recharge its internal batteries making it ideal from a recurring cost standpoint and ease in usability. However, data can only be updated through a computer, and furthermore, does not have a built-in microphone, meaning that end users have no way of including their own new information on the device. Furthermore, it is priced at a point ($45-$65) that is not affordable by the end users (Figure 2.8).

Books of Hope's Speaking Books — At a cost of $10 for a minimum order quantity, Books of Hope's Speaking Books [20] are affordable but is limited to 5-10 minutes of data which cannot be updated. In fact, each new message requires a new Speaking Book to be developed. This causes giving out new information a difficult and costly task.

Given the current contemporary dissemination vehicles and their shortcomings, designing

Figure 2.8: Global Recordings Network's Saber's handcrank allows end users to recharge its internal batteries.

a custom ICT solution from the ground up to face these challenges will yield better results.

2.4 Literacy Bridge's Talking Book

Given the complexities and challenges that need to be overcome listed, in Section 2.2, and the limitations of current solutions from Section 2.3, we argue that developing a custom device is the right avenue to pursue. The Talking Book (Figure 2.9), developed by Literacy Bridge, is a Rural Audio Computer with an audio and tactile interface designed for the sole purpose of disseminating information to illiterate persons in developing regions [21]. Using only buttons, end users navigate through menus to access information and interact with the device. The device prompts the user through spoken instructions which are localized to the region. In January 2009, a pilot study was undertaken in a remote village, Ving Ving, Ghana, where 77% of the adult population had never attended school [5]. Ving Ving's population survives on subsistence agriculture and has no running electricity. In the pilot study, Talking Books were handed out to 20% of Ving Ving's population and were loaded with agricultural information such as instructions about fertilizers, soil preparation, planting and livestock handling. Healthcare practices, educational material, and general stories that had a focus on culture and tradition were also included.

The results of the pilot study, collected in January 2010, showed an increase in crop production by an average of 48% when compared to non end users decrease of 5%. 75% of the end users sold their surplus of crop, and the income from the production was used for health insurance, school fees, home improvements (new roofs), and further investments in their agriculture (seeds, livestock) (Figure 2.10).

The Talking Book can be used for scenarios other than just general agricultural information, such as health practices. For example, nurses in busy rural clinics can record health messages to teach patients about hygiene, infant care, and disease prevention. Treatment specific information can be reviewed at a later time at patients' pace and also be disseminated from a single user that comes back from the clinic to the entire village (Figure 2.11).

For educational purposes, Talking Books can be used to read back to children. This is a teaching practice that is common in developed nations, but impossible for illiterate parents.

Figure 2.9: The Talking Book: A state-of-the-art Rural Audio Computer designed specifically for on-demand information access in developing regions. To access audio content, end users navigate through voice prompts using the small number of input buttons visible on the device.

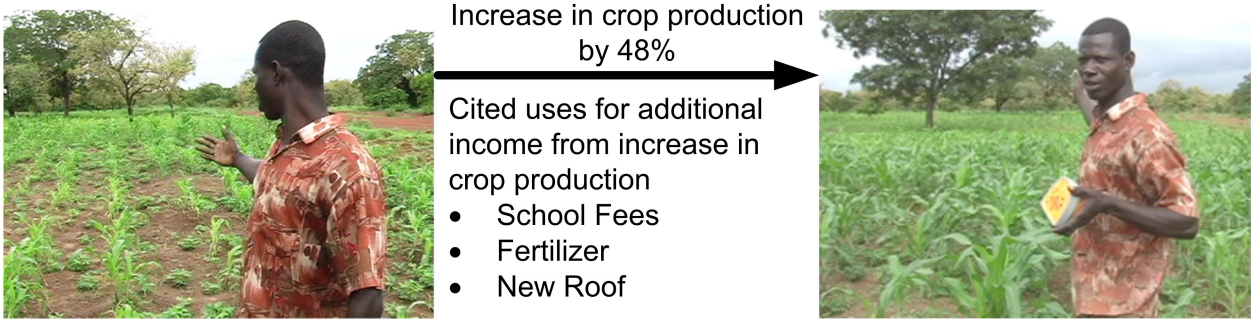

Figure 2.10: Increase in crop production from a user applying information from the Talking Book to his farming practices.

Figure 2.11: A nurse at a rural clinic inspecting the Talking Book for healthcare scenarios.

Figure 2.12: Students in Ghana using the Talking Book in an educational setting.

The Talking Book can also be used as a teaching supplement where questions and answers can be loaded onto the Talking Book and students can be quizzed on material learned (Figure 2.12). This breaks the cycle where illiteracy is commonly passed from parents to children since illiterate parents cannot help their children become literate.

Given the success of Talking Book's pilot study, its chief obstacle to widely distributing it is to make it economically viable. Its current initial cost of $25 - $35 is not affordable by end users, out of the price range that we have previously indicated to be in the $10 - $20 range. The chief contributor to the current cost of the Talking Books is its 152 off-the-shelf components (Figures 2.13 & 2.14). The immediate problem is that the sheer number of COTS components carry significant cost themselves. Purchasing such a large variety of components from different vendors results in every component having some profit margin associated with it. However, there also exist a number of indirect costs associated with the number of components that are not obvious but are in fact, more significant. First, it

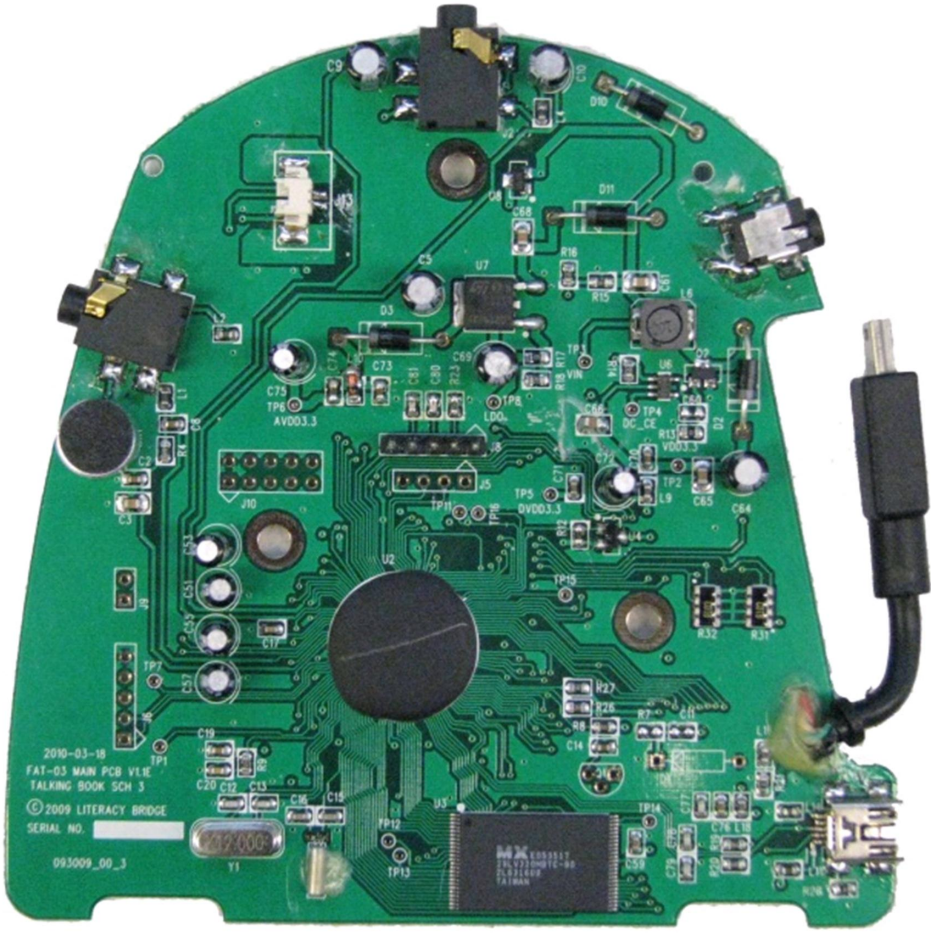

Figure 2.13: Literacy Bridge's Talking Book with its 152 commercial off-the-shelf components (Front Side).

increases the PCB size, the pick and place cost, the probability of failure, and the testing complexity. For instance, the larger device count of the current device requires a PCB area of 120mm X 120mm and a 4 layer board. Finally, the plastic housing, and the shipping costs due to its bulk are increased significantly and all directly translate to a more expensive device. Furthermore, the current Talking Book's high energy consumption requires two D batteries, which impacts its size as well and, again, leads to a costlier device and a form factor that cannot be transported easily by the end user in his/her pocket.

Figure 2.14: Literacy Bridge's Talking Book with its 152 commercial off-the-shelf components (Back Side).

2.5 Implications of Commercial Off-The-Shelf Components Versus Custom Silicon

The main disadvantage of the current COTS component based Talking Book is its cost. As we have outlined in Section 2.4, the Talking Book includes a 16-bit, 96 MHz system on a chip with 256 KB on-chip ROM, a 1 MB NOR Flash chip for program code, and a 256 MB MicroSD card for audio storage. Accounting for 20% of the total bill of materials, the MicroSD card is the most expensive component. The voltage regulator, DCDC converter, speaker amp, and more than 152 other components make up a relatively small portion of the cost but require PCB real estate accounting for 13% of the total cost. We argue that we should merge as many of the COTS components as possible onto a single chip. By doing so, not only can we reduce cost and supply risk, but we also allow for greater PCB flexibility. There are several reasons as to why custom silicon for the Rural Audio Computer would overcome the limitations of the current solutions and are listed as follows.

First, component integration reduces chip count and pin count, both of which reduce cost. Current systems developed from COTS components can consist of as many as 150 discrete components and six major chips [5]. This drives up cost through margin on each individual device and packaging due to high pin-out counts (which can be as high as 50% of the chip cost). By developing custom silicon, we show that the entire system can be integrated in a handful of devices. This saves costs directly through fewer devices and reduced packaging, but also indirectly by allowing a much smaller and simple PCB design (consisting of just 1 or 2 layers). Other indirect benefits also occur, such as lower cost of plastic packaging and reduced shipping costs resulting directly from a smaller PCB footprint.

Second, custom design provides extensive opportunity for a reduction in power draw which reduces TCO by increasing the lifetime on a pair of (disposable) batteries and lowers the recurring cost. Again, indirect benefits of lower power draw exist, such as smaller casing requirements due to smaller or fewer batteries leading to lower fabrication and shipping costs, which accrues to both lower upfront costs and lower recurring costs.

Third, custom integration allows for a design specifically targeted at the developing regions' market. Carbon zinc batteries, for instance, are almost the sole battery option for the

target population. These batteries have significantly different discharge characteristics from alkaline batteries common in developed regions. The choice of battery chemistry, regulator architecture, conversion efficiency, and other factors all play into the design of a low-cost regulator. A custom converter better accommodates carbon-zinc discharge curves. The power supply design is mostly driven by the availability of energy sources in developing regions. From extensive first-hand observations in Ghana, we found that carbon zinc batteries are the most common way of powering electronics (sans mobile phones). The reason is that access to grid electricity to recharge batteries is few and far between, solar cells are uncommon and expensive, and carbon zinc batteries are cheap and have long shelf life. The two most common form factors available are the D and AA sizes. One D size carbon zinc battery holds almost 4X the charge of two AA batteries (8000 mAh vs. 2200 mAh) and costs approximately the same. However, the lower voltage of a single D battery requires that the processor and speaker amplifier operate from a boost converter, which incurs significant energy loss (60% efficient) and also adds cost from additional passive components. Furthermore, the smaller form factor of the AA battery shrinks the overall system size, which reduces cost of the casing and, indirectly, the cost of shipping and distribution. The smaller form factor also allowed easier transport and handling by the end users.

Finally, judiciously designed silicon not only reduces the number of electronic components, but also eliminates other mechanical components. For instance, custom silicon makes capacitive touch sensors economical which eliminates the need for membrane push button switches, which in turn increases reliability and decreases cost. Similarly, the coil necessary for inductive coupling-based, near-field communication (NFC) can be printed on the PCB, eliminating the need for a discrete antenna and additional radio chips. Furthermore, by designing the silicon to be used mostly with capacitors instead of costly inductors allows us to further reduce the cost of the on board passive components.

2.6 Summary

While the statistics and impact of information and literacy on quality of life are well known, the current NGO solutions have significant drawbacks: overwhelming amount of

information that gets lost due to inability to retain that information, and telephone game loss. We show how current commercially available ICT solutions, despite their best efforts, do not adequately solve the unique challenges for information dissemination to populations in rural settings. Some ICT solutions are out of the price range of the end users, while other ICT solutions do not take into account the lack of infrastructure, content delivery and replayability. Some ICT solutions while within the price range, are not robust given their electro-mechanical actions. Given the complexities needed to be overcome in this unique situation, we see that custom design is a logical choice to cut the Gordian Knot that ties cost, power, connectivity, usability, robustness, and illiteracy. However, as the Talking Book has shown, custom design even at the board level has been proven insufficient to put the device at a cost point to be purchased by end users. Hence, custom silicon is required.

CHAPTER 3

Overview of Literacy in Technology (LIT): A Low-Cost Custom Silicon Rural Audio Computer for Information Dissemination among Illiterate People Groups

3.1 Motivation

In Chapter 2, we discussed how current ICT solutions do not adequately overcome the challenges of disseminating information among illiterate people groups. We argued that the unique challenges — cost, power, connectivity, and robustness can be overcome through the design of custom silicon. As previously mentioned, the chief challenge to disseminating the ICT devices among this list is cost due to the low purchasing power of the end users that typically subsist on \$1-\$2 per day. Based on our analysis of electronic devices that are frequently purchased in the market places such as radios and cassette recorders as well as through discussions with various NGOs, we believe a total device cost must lie between \$10-\$15, which means that all the electronics, including PCB, speakers, microphone, radio, buttons, etc. must remain below \$6.

In order to meet this \$6 cost target, this thesis proposes an SOC which forms the heart of an audio computer for information dissemination among illiterate people groups. This people group has a currently untapped market size of \$10B [22]. In this chapter, we give a high-level overview of the features and components of the SOC, referred to as the Literacy Information Technology (LIT) chip. We also discuss at a high-level how these features help to

meet the $6 cost target as well as the other challenges of power, connectivity, and robustness in a novel way. We present each of the key features of the LIT chip in more detail in the following chapters.

3.2 Overview of LIT Chip and Approach to Reduce Cost

With the LIT SOC at its center, the audio computer allows users to navigate through menus using buttons enabling them to access information which is played back to them through a speaker. The audio computer's information can be updated either through an on-board FM chip that the SOC communicates with, through a microphone that is connected to an on-chip ADC, or through peer-to-peer distribution via its near-field radio link. Using audio feedback, the user can configure and operate the device in a similar way that made the Talking Book's pilot study a success.

In figure 3.1, we show the overall architecture and its subcomponents. LIT was implemented in 0.18um CMOS, measures 3.57mmX6.46mm (23.06mm^2), has 8 million transistors, 265K gates, with an expected cost $<$ \$1 in moderate volumes (Figure 3.1). LIT chip tackles cost in a number of ways. For example, a novel memory architecture with 128kB 4-way True LRU cache is directly backed by NAND Flash and removes the need for costly DRAM or NOR Flash commonly found in micro-controllers and processors.

LIT also integrates all analog components in the system on-chip, including microphone amplifier and biasing circuits, an Analog-to-Digital Converter (ADC) [23], a class-D amplifier, 4 Low-Dropout Regulators (LDO)s, 2 voltage references, and a Hybrid Switch Capacitor regulator for DC-DC conversion. The high level of integration onto a single die reduces the total number of chips in the system to only 3: the LIT chip, a NAND Flash chip, and a radio chip. This saves cost and PCB size, which in turn reduces the plastic and handling costs (the previously mentioned indirect costs) of the final ICT solution.

There are current solutions that have a high level of integration [24] [25] [26] that meet some of the capabilities that fit the current ICT goals of information dissemination. Also,

high levels of integration of SOC reuse has been known to save design complexity, product cycle time, and production cost [27] [28] [29] [30] [31]. However, LIT is the first custom chip designed specifically for disseminating information among low-income populations. Its high-level of integration results in a 9X reduction in components when compared to Talking Books [21] and other current ICT solutions available [13] [20] [32] which reduces both direct and indirect costs.

3.2.1 Discrete Components

Even after integrating as many components on the LIT SOC as possible, there are still remaining discrete components. To reduce the cost of these components, we have also made design choices that directly reduce their cost. For example, instead of using a Step-Up Boost Converter with an expensive inductor for voltage boosting, we developed a new fine grain Hybrid Switch Capacitor Network which uses only capacitors. Since capacitors are one order of magnitude cheaper than inductors, this reduces overall system cost. Furthermore, peer-to-peer information transfer is achieved through a coil that is directly traced on the PCB for a near-field inductive link for close range communication between devices instead of costly USB connectors/wires/antennas. This design also increases the robustness of the system by not requiring wires or connectors. In addition, 10 Capacitance-to-Digital Converters (CDCs) [33] were integrated on-chip to allow push-membrane buttons, which incur significant cost, to be replaced with capacitive sensors, also traced directly on the PCB. This overloading of PCB traces is unique in the field of ICT devices and results in a cheaper, more robust overall solution.

A side effect of the effort to reduce the number of components also results in a more robust device. Replacing the Talking Book's USB dongle (Figure 2.9) with the coil for peer-to-peer transfer reduces the number of physical open ports, and the capacitive sensors on the board results in a longer lifetime of the buttons since push-button membranes will deteriorate with use (Figure 3.2). While robustness may not be as large a concern as cost, prior ICT devices have shown that end users are prone to disassembling the device and leaving it in the mud/rain indicating that end users may not be as careful with the device as typical users

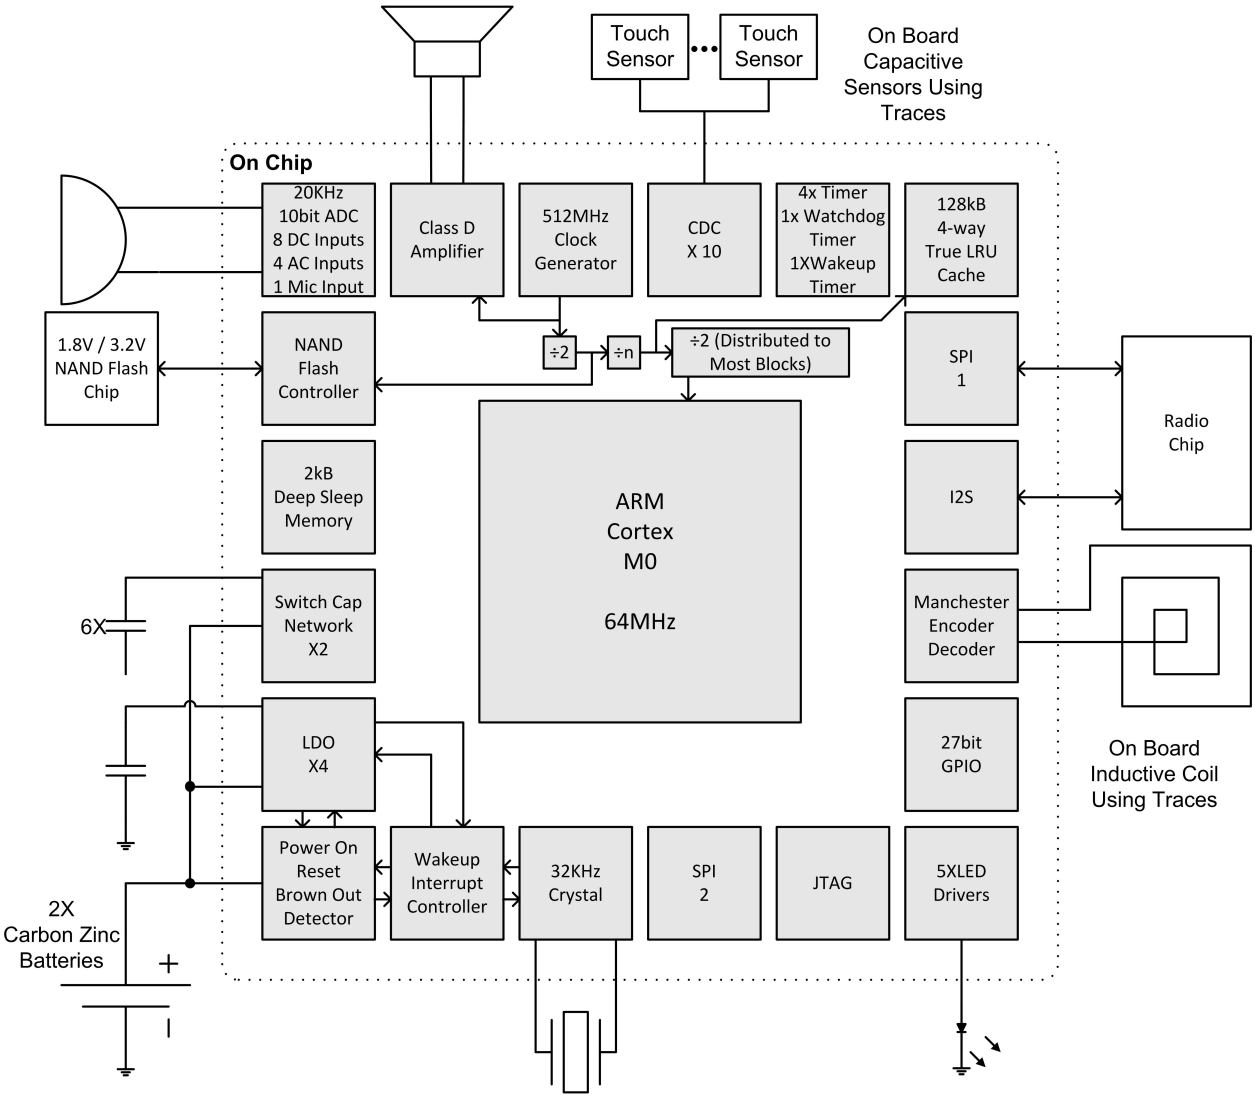

Figure 3.1: LIT's System Level Diagram Showing a High Level Integration of On-Chip Components Resulting in a Low On-Board Component Count.

that are more accustomed to electronic technology (Figure 3.3).

3.2.2 Power Consumption

Lowering active and sleep energy consumptions are important goals for LIT since this reduces the end users' recurring costs, and many of LIT's design choices were made with lowering energy consumption in mind. Clock gating, power gating, and clock speed tuning, depending on workload and module activity, are software configurable and work in unison to lower energy consumption across the board. Furthermore, most of LIT's power hungry

Figure 3.2: Broken OLPC Push-Membrane Keyboard.

Figure 3.3: Talking Book left outside in mud overnight by end users.

components can be power gated individually depending on their use, thereby only using the minimal amount of energy required to do exactly what the device is doing at any point in time. LIT's low energy consumption allows it to be operated with only two AA batteries for up to two years in Deep Sleep mode, which improves its form factor to a point that makes it more easily transported by end users and lowers the recurring costs of purchasing new batteries for the device.

LIT's processor is an ARM Cortex-M0 which was chosen to balance power consumption with performance. The ARM Cortex-M0 uses only ∼7000 gates from ARM's TSMC180 Standard Library, thereby requiring only modest chip area and contributes to the LIT chip's area by a minimal amount. The ARM Cortex-M0 has 32bit data operations and is powerful enough to decompress Speex audio, constant pitch speed shift audio using Synchronized Overlap-Add Fixed Synthesis (SOLAFS), and manage system overhead simultaneously at

a clock speed of 64 MHz. Since the Cortex-M0 uses thumb code, consisting of 16-bit instructions as opposed to 32-bit instructions, it has high code density and minimizes system memory size. Its current draw of only 5mA during operation limits the recurring cost due to energy consumption.

3.3 Communication Capabilities

In Chapter 2 we showed that one of the drawbacks of current ICT devices is with their communication capabilities. LIT overcomes this through the integration of multiple communication modes with the inclusion of FM Radio, Peer-to-Peer communication, and mobile phone communications capabilities while maintaining a low cost. For wide-area dissemination, LIT uses an off-chip Silicon Labs FM radio receiver chip. A custom inductive link communication interface supports peer-to-peer communications, and connectivity over a mobile phone is possible by transferring data through the audio headset port [34]. We discuss each communications modalities briefly below, with the exception of the custom inductive link communication interface which is discussed in Chapter 8.

FM Radio — In order to receive wide-area programming, we use the Si4705 FM Radio receiver chip from Silicon Labs [35]. We chose an external chip rather than attempting to design and integrate the equivalent functionality onto our own silicon to mitigate numerous incompatibilities, risks, and challenges involved with implementing and RF receiver on silicon. Many parts of the radio chip are large analog components that tend to use older semiconductor process technologies that are incompatible with the rest of our custom silicon and would increase its size, cost, and power draw. Furthermore, while the cost of the Si4705 is significant, it adds substantial value to the platform as the chip can receive both FM audio and RDS concurrently. The only additional capability that was required to add to our own chip with regards to FM is the 32KHz timer that allows us to schedule recordings at future times. That way, the user need not wait until a particular program is broadcast, and can instead have content downloaded automatically. While other ICT solutions offer this capability as well, there were design decisions that allowed us to overload the use of the 32KHz timer used by the radio thanks to the co-dependency of having as many components

on a single die. This allowed us to reuse the crystal for accurate time keeping and low clock frequency monitoring of events which resulted in a lower energy consumption. Furthermore, the 32KHz accurate timer allows us a wake up and "Tivo" recordings from the radio, a capability that is not available on any of the ICT solutions.

Mobile Phone — To support communications with a remote server over the GSM network, we added a Universal Asynchronous Receiver/Transmitter (UART) controller to our chip. The UART controller is small (less than 1% of the chip area) and adds only two additional pins (RX and TX). The rest of the processing is done on the Cortex-M0. In addition, an audio crossover cable can be used to connect the device's earphone port to a mobile phone's headset port. None of the current ICT devices in this application space support this solution and this allows LIT to have more flexibility in its communication capabilities.

3.4 LIT's Novel Solutions

This thesis presents a number of solutions that were used in the LIT chip implementation and were specifically developed to address the needs of ICT in developing world applications. Since the objective was to optimize the overall cost, robustness, power, and connectivity of the final solution, the novelty of the approaches vary from low level circuit techniques to system level. To help clarify the novel aspects of the presented work, we use three novelty categories as listed below. The table in Figure 3.1 indicates for each solution presented in this work which novelty categories we believe it belongs based on our understanding of the prior art.

Innovation categories:

- Novel Techniques: Solutions in this category are specific techniques that novel in and of themselves. They improve upon previous techniques in terms of how well they meet their objectives and they were not previously proposed in the literature. Examples of novel techniques are circuit methods, such as LIT's DC-DC step-up converter. The LIT DC-DC converter uses a circuit structure that was not previously proposed and improves one of its metrics, specifically, the voltage tuning granularity of the converter.

- Novel System Solutions: Solutions in this category use existing and known techniques, but configures them at the system level in a novel manner. Hence, the novelty does not lie in the techniques themselves, but how the techniques are composed to create a larger system (in this case, the LIT system) to accomplish a particular goal. For instance, the memory hierarchy used in LIT uses a number of techniques that are well known in the literatures, including a low power microcontroller, a cache with hardware based miss detection, software based cache miss handling, etc. However, the particular combination of (a) a low power and low cost commercial core without built in memory management capabilities, (b) a large 128kB cache with a hardware based cache miss detection and software based cache filling, and (c) a large NAND Flash based memory storage, is unique, not commercially available, and is key to enabling a large code base and memory space, while maintaining programming ease with extremely low hardware cost.

- Novel Application Specific Solutions: Solutions in this category are methods that have been previously proposed in a different applications space but were found to have a unique benefit for the application at hand and were not previously used in this application space. Likely it was used in the original application space for different reasons that what makes the method applicable to the current application space. The contribution in this case is the recognition that this method can fill a particular need in a unique and beneficial way in the studied application space: ICT devices for developing world applications. An example of this type of novelty in the LIT chip is the use of capacitive sensing switches. While capacitive sensing is well known as a technique for data input (from microwave ovens to cell phones), the recognition that they hold particular advantages in terms of robustness and cost for the LIT application space is new and their use in this application space was not previously reported. In fact, the association of capacitive switches with high-end products in traditional markets makes it counter intuitive that it would actually lower cost instead of increase cost in our LIT application. However, since the LIT chip is a VLSI design, the fairly sophisticated capacitive sensing circuits could be included on the chip without increasing die size

substantially and therefore comes essentially for free. Since it removes the need for membrane switches, which are actually fairly expensive at approx. $1, their use actually reduces overall system cost substantially.

Table 3.1: LIT's Novel Design Decisions.

LIT's Design Decisions	Novel Technique Solution	Novel System Solution	Novel Application Specific Solution
Low-Cost Hybrid Switch Cap Network DC-DC Converter	**X**		**X**
Parallel sleep & active LDOs to to quickly regulate and switch between operational modes			**X**
POR/BOD able to overcome high Carbon-Zinc battery hysteresis	**X**		**X**
Configurable Pad Voltage using on-chip voltage regulation		**X**	**X**
Touch Sensors & Communication Link using only on-chip components & PCB based traces			**X**
128kB true LRU cache directly backed by NAND Flash to eliminate code NOR Flash			**X**
High level of system integration with low number of passives & discrete chips			**X**

In the following chapters, we discuss the different unique features of the LIT chip in more detail and discuss their novel aspects. Specifically, Chapter 4 discusses LIT's power management techniques and design strategies to provide a low cost robust power solution that provides power from a single Carbon-Zinc battery source. Chapter 5 highlights LIT's novel Hybrid Switch Capacitor Network, designed to power off-chip components such as a NAND Flash chip or Silicon Labs radio. Chapter 6 explores the novel memory hierarchy used in LIT that uses a co-designed software managed cache with hardware cache miss handling and removes a costly DRAM or NOR Flash. Chapter 7 shows how LIT maintains robustness through the design of a novel Power-on-Reset / Brown-out-Detector with Lock-Off for Carbon-Zinc batteries. Finally, Chapter 8 discusses how LIT uniquely lowers cost

and increases robustness through the overloading of PCB traces for use in a human input interface and data transfer communication methods that do not require additional dongles, physical open ports or wires. Given the design decisions and constraints that LIT has, we show LIT's results in Chapter 9 and conclude with Chapter 10.

3.5 Summary

To summarize, we provided a brief, high level overview of the functionality of the LIT chip and its unique features. We also showed how careful design of custom silicon can overcome the unique challenges — cost, power, connectivity, and robustness for the purpose of disseminating information among illiterate people groups that improves upon previous ICT designs. We achieve this primarily through, a high-level of integration, design decisions and novel structures that reduce the cost of remaining discrete components, and low power design solutions. In the following chapters, we discuss each of the key components of the LIT chip in more detail that arises out of this unique challenge, application space, and high-level of integration.

CHAPTER 4

On-Chip Integrated Power Management, Power Operational Modes and Wakeup Interrupt Controller

4.1 Introduction

Chapter 3 showed how a high level of integration can result in a low cost ICT device that meets the $6 cost target that can be bought by the end users who subsist on $1 to $2 per day. In order to satisfy this constraint, LIT's power management system is generated entirely on-chip derived off a single Carbon-Zinc battery while only using cheap passive components such as capacitors off-chip. LIT lowers its operational cost through the minimization of energy consumption through multiple methods. 1) The LDOs can be bypassed and directly connected to the Carbon-Zinc battery during the battery's voltage degradation to eke out every last bit of charge. 2) Having multiple operational modes in order to reduce energy consumption all supported by its Wakeup Interrupt Controller. 3) Variable clock frequencies and module power gating. LIT also reduces cost through being sensitive to market prices of the of external NAND Flash Chip through its Board Level Configurable Pads where LIT can communicate to either a 1.8V or 3.2V NAND Flash Chip powered by either it's on-chip Dirty LDO or Hybrid Switch Capacitor Network. LIT handles power in a novel systems and application specific method not attempted by commercially available solutions or current ICT solutions that allows it to achieve both a low initial and recurring cost, and we expand upon them in this chapter.

4.2 LIT's On-Chip Integrated Power Management

4.2.1 Overview

In total, LIT has 4 LDOs [36] [37] [38] [39] [40] (Figure 4.9) with their associated VREFs [41] and IREFs [42] and a Voltage Doubler (Figure 4.1). LIT's 4 LDOs are used to generate on and off-chip 1.8V and off-chip 0.95V voltage domains. LIT also has a single Voltage Doubler used to generate a 3.2V voltage domain that is used to power several Thick-Oxide devices for energy efficient circuits and for LIT's off-chip communication.

4.2.2 LIT's 4 On-Chip LDOs and Unique Flexibility

LIT has two major LDOs that are used to power its on-chip 1.8V power domain, its Active and Deep Sleep LDO. These two LDOs operate in parallel during LIT's Active mode, when LIT is doing most of its heavy computation and processing. When the Active LDO is enabled, LIT can draw more current and consume more energy for its workload (mW range). However, when LIT switches to Deep Sleep mode, LIT isn't doing much and therefore goes into a low energy consumption mode, LIT's energy consumption decreases by 1000X (uW range), so the Active LDO is then power gated and the Deep Sleep LDO takes over regulating the on-chip 1.8V power domain. This allows LIT to consume less energy since the Deep Sleep LDO is more efficient, but cannot sustain as high a current draw as the Active LDO. During Active mode, while the Active LDO is enabled, the Deep Sleep LDO is also enabled in parallel, but overpowered by the active LDO. This allows for an easier handoff during switching between Active and Deep Sleep modes where we can easily enable or disable the Active LDO. Having dedicated LDOs for LIT's Active and Deep Sleep modes allows us to achieve high energy efficiency and to maintain a low Deep Sleep mode energy consumption that provides LIT up to 2 years of lifetime. Although parallel Active and Deep Sleep LDOs techniques are known [43], typical microcontrollers [24] [25] [26] only have at most, one LDO on-chip. Therefore, in order to have two separate LDOs for power saving modes, a separate single channel LDO chip with adjustable outputs would be required [44], but this is counter to LIT's design goals of minimizing the number of off-chip components, and so LIT integrates

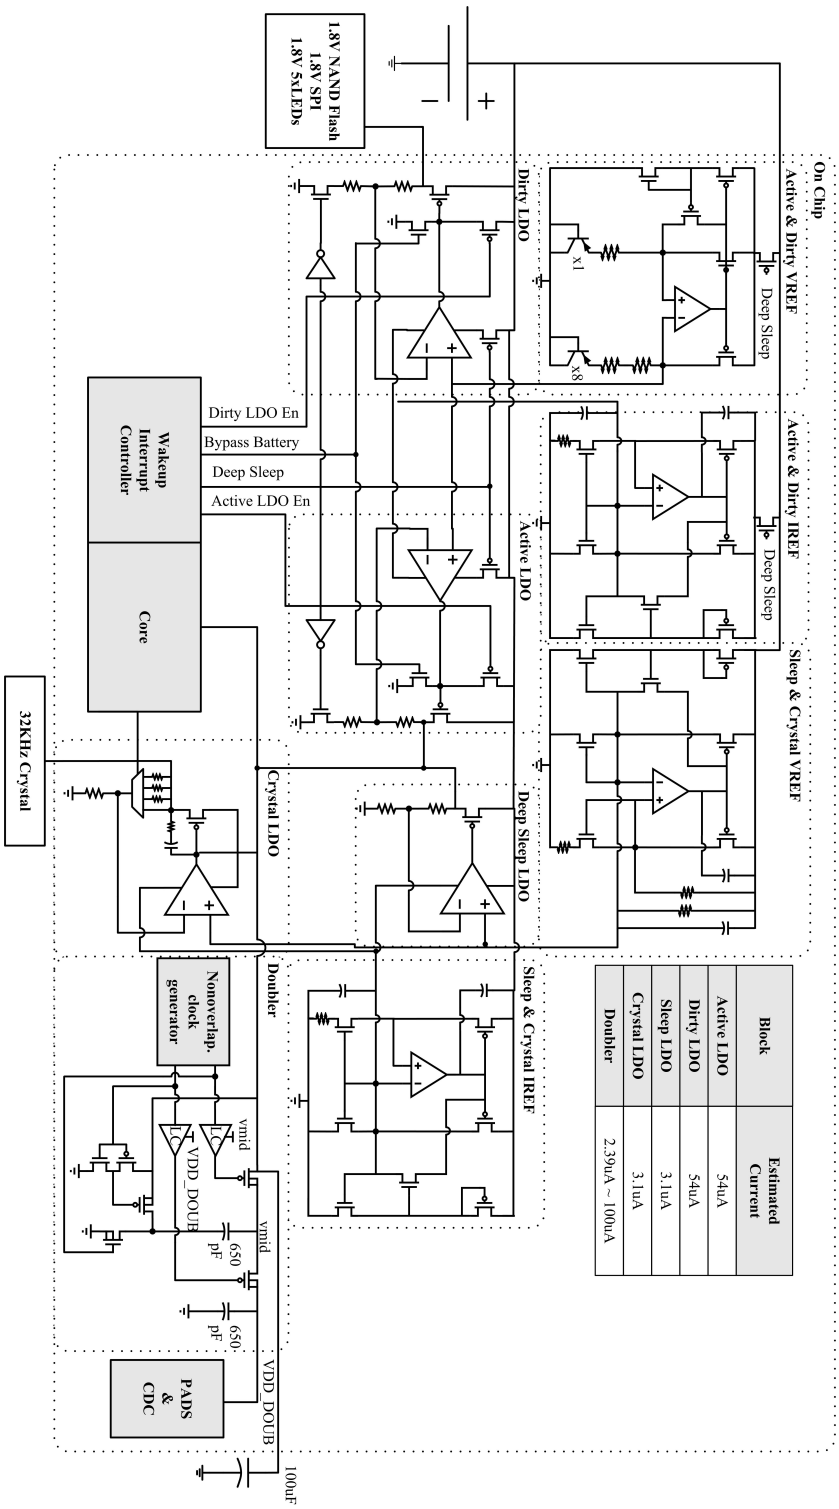

Figure 4.1: LIT's power diagram showing LDOs, and on chip Voltage Doubler with their associated VREFs and IREFs and silicon measured linearity.

all LDOs on-chip.

LIT's third LDO, the Dirty LDO, is used power off-chip 1.8V power domains and components such as a 1.8V NAND Flash Chip. Since LIT's memory hierarchy requires us to have an on-board NAND Flash chip that is cheaper than other solutions (Chapter 6), LIT maintains some flexibility in that it is capable of communicating with either a 1.8V NAND Flash chip or a 3.2V NAND Flash chip. The compatibility and flexibility is important because while the 1.8V NAND Flash chip consumes lower energy, resulting in a lower recurring cost for the end user when compared to the 3.2V NAND Flash chip, the NAND Flash chip market is volatile, and can result in a 1.8V NAND Flash chip whose initial cost is too expensive, making LIT unfeasible. In order to provide a solution to power two alternatives, LIT can power the 3.2V NAND Flash chip with its Hybrid Switch Cap Network (Chapter 5), while powering the 1.8V NAND Flash uses LIT's Dirty LDO. In the event that a 3.2V NAND Flash chip is used, The Dirty LDO can be entirely disabled in order to minimize energy consumption. By having the Dirty LDO, this flexibility allows LIT to be sensitive to market volatility for its dependencies on external components allowing us have the ability to choose the lowest cost point. LIT is unique in that it is flexible in its choice of possible external components with its entirely on-chip power management system when compared to the previous ICT solutions 2.

Finally, LIT's 4th LDO, the Crystal LDO is used to regulate an off-chip 32KHz crystal's power domain. LIT uses an off-chip 32KHz crystal to tune the radio chip and to keep real time using an on-chip timer. This timer allows LIT to "Tivo" content from the radio chip so that content distributers can broadcast firmware updates through FM RDS at predetermined times. Additionally, users can set predefined times to wake LIT up and automatically record new content from the radio that can be listened to later. This is important to end users since government agencies freely distribute content on national radios so that users can record and review the information later at their own pace. In order to power the 32KHz 830nW off-chip crystal, LIT's 4th LDO, the Crystal LDO (Figure 4.1), constantly regulates its 0.95V supply, and is always enabled, like the Deep Sleep LDO. The 32KHz clock generated from the crystal needs to constantly be enabled since LIT's triggers to exit Deep Sleep are clocked using this clock. Furthermore, LIT's Wakeup Interrupt Controller used to monitor these

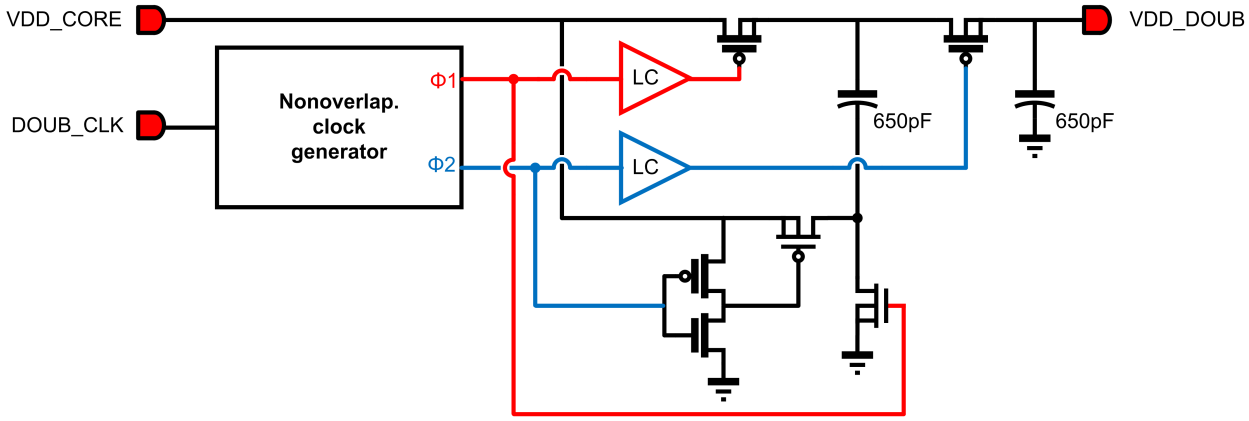

Figure 4.2: LIT's Voltage Doubler's Circuit Level Diagram.

triggers operate in the 32KHz clock domain and therefore must always be available.

4.2.3 Voltage Doubler

LIT has a few components on-chip that require a 3.2V power supply such as LIT's Pads' Dirty VDD, for off-chip communication, and the CDCs (Chapter 8), that both operate at 3.2V. In order to supply the 3.2V, LIT has an on-chip Voltage Doubler (Figure 4.1) (Figure 4.2) (Figure 4.3) that is similarly derived from the Core 1.8V like the Crystal LDO. The Voltage Doubler is also required to stay alive during Deep Sleep mode in order to power the CDC that monitors the PCB touch sensors, thus, we need to pay attention to its energy consumption. The Voltage Doubler has a configurable clock speed that can either take the Core Clock speed that runs in the MHz range during Active mode, or can switch over to its own internal clock generator that operates in the 100s of KHz range, for example during Deep Sleep mode or if there is not much off-chip communication. This allows the Voltage Doubler to operate in an energy efficient manner and help achieve low energy consumption during Deep Sleep mode. While the system's technique of charge pumping a regulated load for internal power delivery is known even in commercial chips [45], commercial chips don't nearly have the same amount of functionality that LIT provides, and it is priced at a point [46] that is the total electronics cost for LIT.

LIT's On-Chip Integrated Power Management allows LIT to achieve its high level of integration since it only requires off-chip capacitors. This reduces its cost while being able

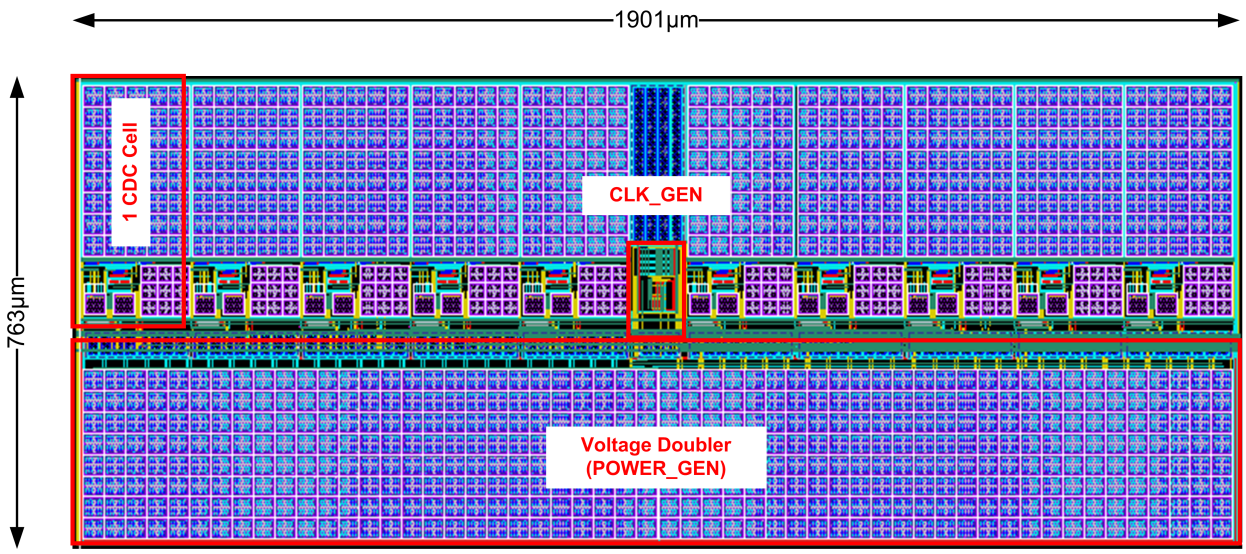

Figure 4.3: LIT's Voltage Doubler's Layout Measuring 1.9mm X 0.35mm Occupying 3% of LIT's Area.

to maintain multiple power domains all generated from a single chip solution. LIT's Power Management also allows it to be to choose between a 1.8V or 3.2V NAND Flash chip making it flexible to the market volatility of its prices. These strategies improve LIT when compared to previous ICT designs, and while most commercial systems have integrated power solutions, none boast as many and as flexible as LIT does, allowing it to both achieve its low initial cost, and low energy consumption therefore minimizing its recurring cost to the end user.

4.3 Minimizing LIT's Energy Consumption

As we have mentioned numerous times, LIT's main goal is to lower its cost, but we should not only look at LIT's initial cost, but also look at its recurring cost. In order to reduce its recurring cost, LIT must minimize its energy consumption. We go into detail LIT's energy minimization techniques through its Carbon-Zinc battery direct connect, its power operational modes, wakeup interrupt controller and its variable clock frequencies and module power gating in the following subsections.

4.3.1 LIT's Carbon-Zinc Battery Direct Connect

LIT was designed to maximize lifetime when powered by Carbon-Zinc batteries [47], prevalent in developing regions due to their superiority over Alkaline batteries in terms of longer shelf life and lower acquisition cost [48]. Carbon-Zinc batteries tend to degrade slower over a larger voltage range unlike Alkaline batteries which have a steep cutoff voltage at 2.2V (Figure 4.4). In order to extract maximum energy, it is critical to extract energy from the battery over the entire voltage degradation range [49]. Therefore, LIT uses a two phase regulation scheme: 1) At high battery voltages, LIT's LDOs regulate the battery to the 1.8V core supply with a low drop out of 140mV; 2) Below 1.9V (ie: 0.85V per battery), we by-pass the LDOs and directly connect the battery reducing our cut-off battery voltage to 1.7V. This is achieved through monitoring the battery voltage with the on-chip ADC. Software will periodically use the ADC to check the battery voltage to determine if it should directly connect to the battery. This allows us to extract 76% of total charge vs. 46% of total charge at 2.2V (Figure 4.4) thereby extending the lifetime of the device and lowering recurring cost. Three of LIT's LDOs, the Active, Deep Sleep and Dirty LDOs use the two phase where we directly connect the battery to the Core's 1.8V when the battery's voltage drops below 1.9V, but the Crystal LDO is unaffected by the battery bypass since its power supply comes from the Core 1.8V supply instead of the battery. This allows LIT to have a more energy efficient Crystal LDO since it has a much smaller input voltage range when compared to the battery voltage.

4.3.2 LIT's Power Operational Modes and Wakeup Interrupt Controller

In order to fully exploit LIT's multiple LDOs, LIT was designed to have 3 operational modes: Active mode, Standby mode, and Deep Sleep mode in order to reduce energy consumption. LIT needs some way of coordinating all its operational modes (Active, Standby, Deep Sleep), managing the power supplies, (LDOs, Carbon-Zinc Battery Bypass), power-gating the cache and clock generation, and monitoring of the Deep Sleep event triggers while consuming as little energy as possible since it would be required even during Deep Sleep

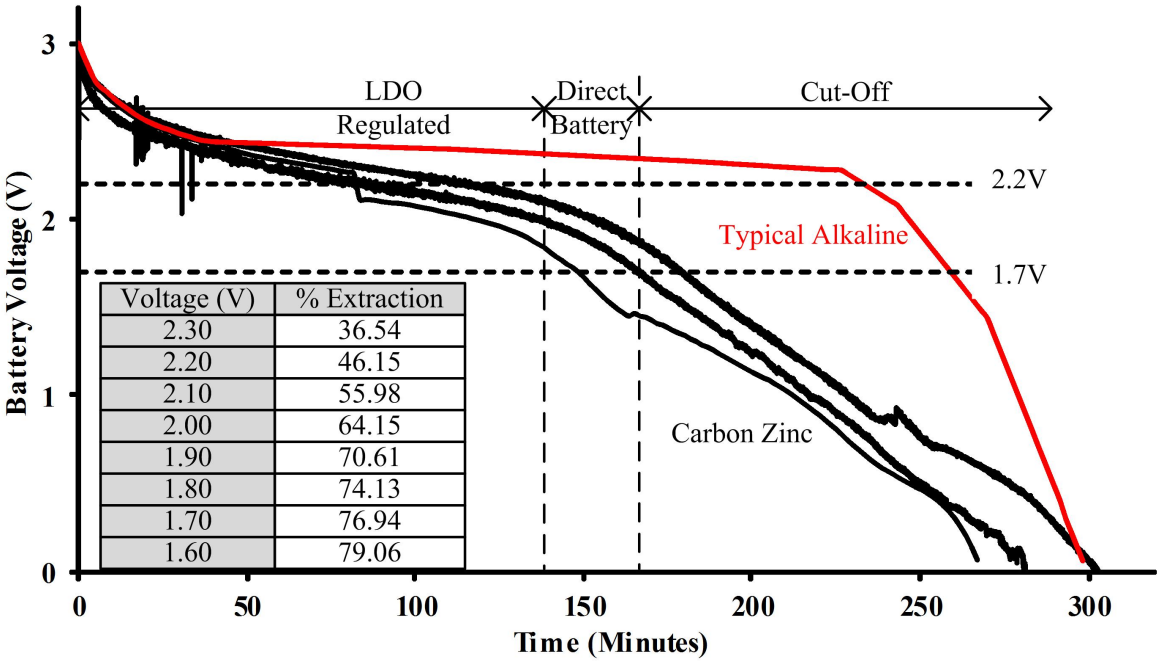

Figure 4.4: Measurements of Ghana's Carbon-Zinc batteries versus typical Alkaline battery. By extending operational voltage from 2.2V to 1.7V, 30% more energy is extracted.

mode [50] [51] [52] [53]. Thus, we created a Dynamic Power Management system: LIT's Wakeup Interrupt Controller (WIC) (Figure 4.5) which achieves all the above. Figure 4.6 shows how the WIC controls the LDOs, Clocks, and Caches when it goes in and out of Active and Deep Sleep. The WIC allows LIT to achieve its minimum energy consumption in order to lower its recurring cost to the end users.

The WIC has a low gate count to minimize energy consumption, uses only 294 gates, 103 of which are Flip-Flops. 32 of those Flip-Flops, which maintain state, operate on the 32KHz

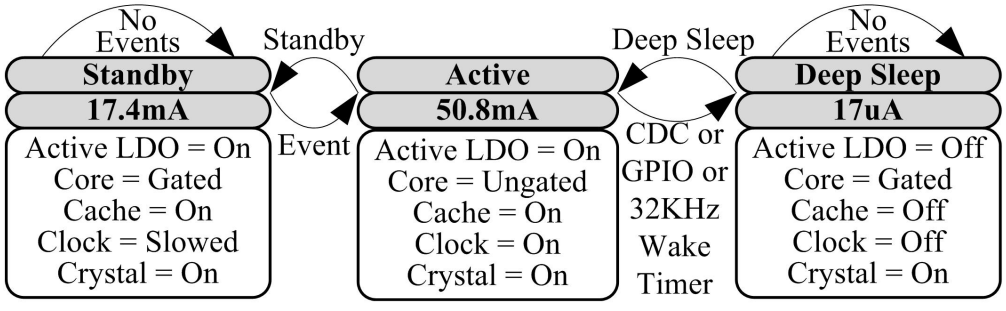

Figure 4.5: LIT's Wakeup Interrupt Controller state diagram.

clock domain, ensuring that it consumes very little energy during Deep Sleep operations while the other Flip-Flops, used to store configuration bits, have its clock generation disabled.

During Active mode, while LIT is in operation, the Active LDO is operating as described in Section 4.2. When software detects that the Battery Voltage has drooped too low through the on-chip ADC, software informs the WIC and the WIC will inform the LDOs to directly connect the Core Voltage to the Battery. The WIC also controls whether the Dirty LDO is enabled or not and also power gates Active and Dirty LDO and their associated VREFs and IREFs during Deep Sleep mode.

Since LIT's operation relies on user input that could range in the seconds, LIT has an operational mode between Active and Deep Sleep mode: Standby mode. In Standby mode, the Cortex M0 is clock-gated and the main system clock is slowed until an event occurs, and draws a measured 15.3mA. The Standby mode allows for lower energy consumption between the user doing something and the software deciding to shut down and enter Deep Sleep mode.

Lastly, LIT has a Deep Sleep mode where the Active LDO is disabled, the on-chip clock generators are disabled and most modules are power gated. During Deep Sleep mode, the WIC monitors the event triggers, such as the human input interface (CDCs), GPIO inputs, or the 32KHz Timer. During Active mode, the CDC Controller and GPIOs operate on the active Core Clock domain (MHz), but during Deep Sleep mode, they are switched over to the 32KHz clock domain in order to allow the WIC to disable the clock generator and consume less energy. Since the WIC operates on the 32KHz clock domain, the initial cost of having an operational clock is amortized amongst modules that are left awake in Deep Sleep mode. LIT's Deep Sleep mode allows LIT to last $\sim$ 2 years. This low energy consumption lowers its recurring cost to the end user, making LIT more affordable. The WIC is also in charge of controlling the power gating of the cache and the clock enabling. The different cache power gating configurations are set as configuration bits in the WIC during Active mode and the WIC will power gate the entire cache or leave subsections alive as described in Chapter 6. In Deep Sleep mode, the clock generator is halted, the Active LDO, potentially the cache, and unused blocks are power gated. The WIC which monitors the CDC controller, 32KHz Timer, or GPIO will wake up the Cortex M0 in an event and draws a measured 17μA

(32KHz WakeTimer Monitoring) to 27μA (All Event Monitoring) that results in 2 years of Deep Sleep lifetime (Table 4.1).

Table 4.1: LIT's Deep Sleep mode current draw with event monitoring options selected.

Deep Sleep Event Monitors	**Measured System Current**
32KHz Timer	17μA
32KHz Timer & GPIO	22μA
32KHz Timer & GPIO & 10XCDC	27μA

The WIC also has a wakeup protocol when emerging from Deep Sleep. Since many of our modules are power gated during Deep Sleep, immediately enabling all our modules would result in a voltage droop that may lead to undefined behavior. LIT's wakeup protocol results in more robust, but slower wake up time where we first enable the Active LDO, wait for a configurable number of 32KHz clock cycles, then enable the clock thus giving the Active LDO ample time to be operational. We also only restore power to the cache after a 32KHz clock cycle and wake up the Cortex-M0 another 32KHz clock cycle after waking the Cortex-M0. This ensures that we do not have undefined behavior due to too much current being drawn simultaneously and results in a more robust system (Figure 4.6).

4.3.3 LIT's Variable Clock Frequency and Module Power Gating

LIT's workload can vary greatly depending on its usage, and we can alter its clock frequency, power gate blocks depending on usage. This allows us to control energy consumption through software to be at an energy efficient point for its current workload during Active mode. For example, if audio is only being played or recorded, LIT can fill the audio buffer completely and then slow down the core clock while audio is being played. The other side of the coin is when LIT requires all its processing power at 64MHz. End users have found time shifting of the information to be useful by having the information read back at a faster/slower speed to help. However, this requires use of the synchronized overlap-add method (SOLAFS) which requires a high amount of processing power which LIT achieves with its 64MHz clock speed. Thanks to LIT's workload aware variable clock frequency and ability to power gate blocks, LIT draws a measured current of 17.4mA with only the Cortex M0 running at 4MHz and 50.8mA at 64MHz (Figure 4.7).

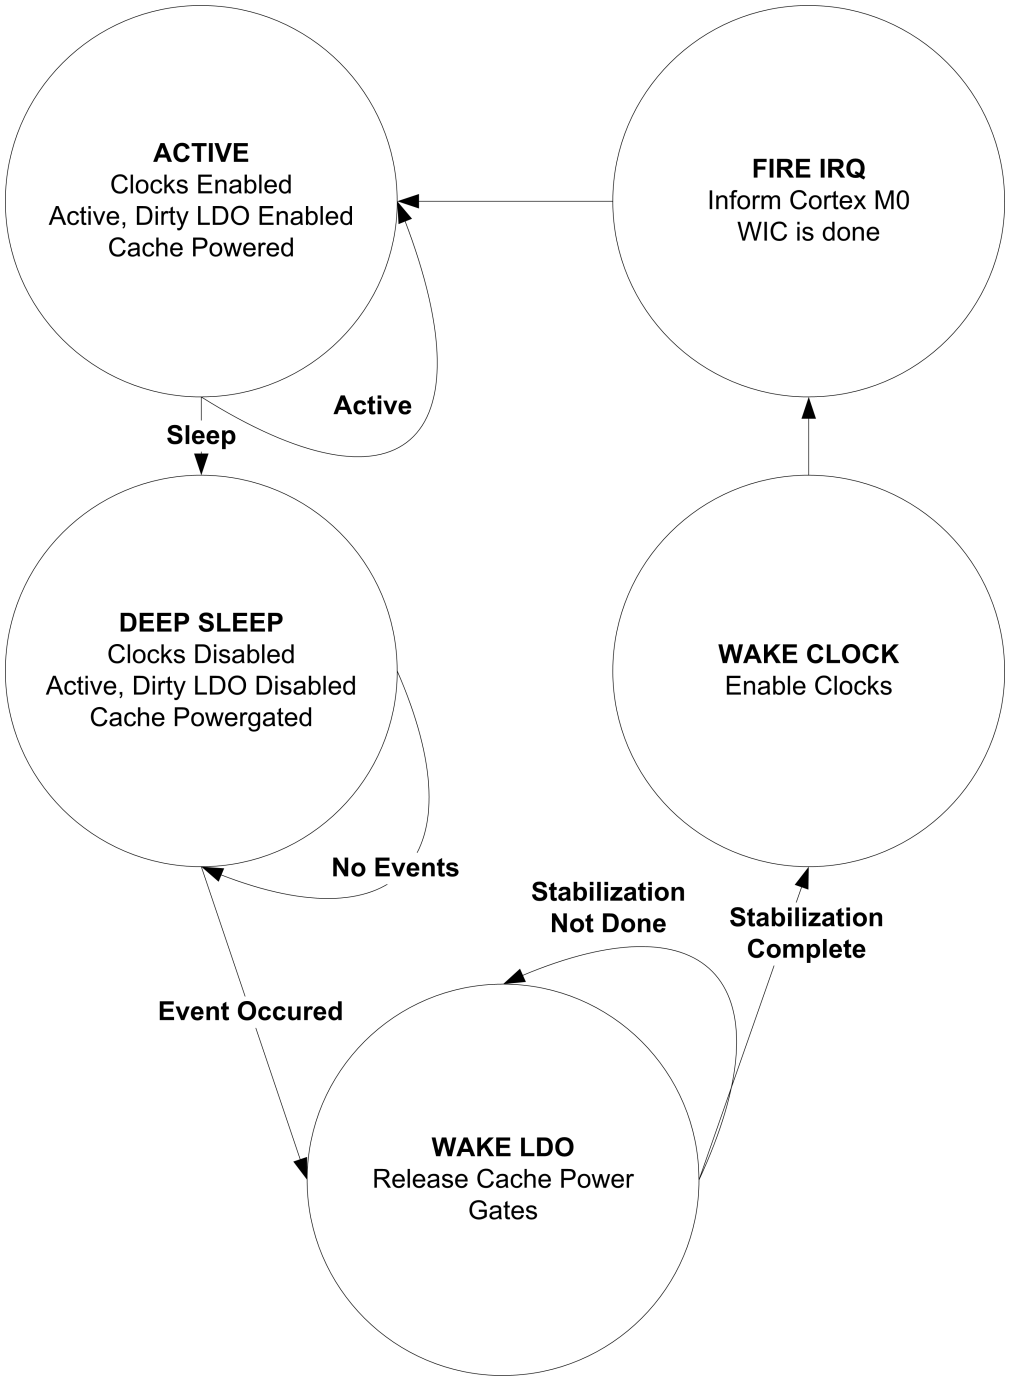

Figure 4.6: LIT's Wakeup Interrupt Controller state diagram.

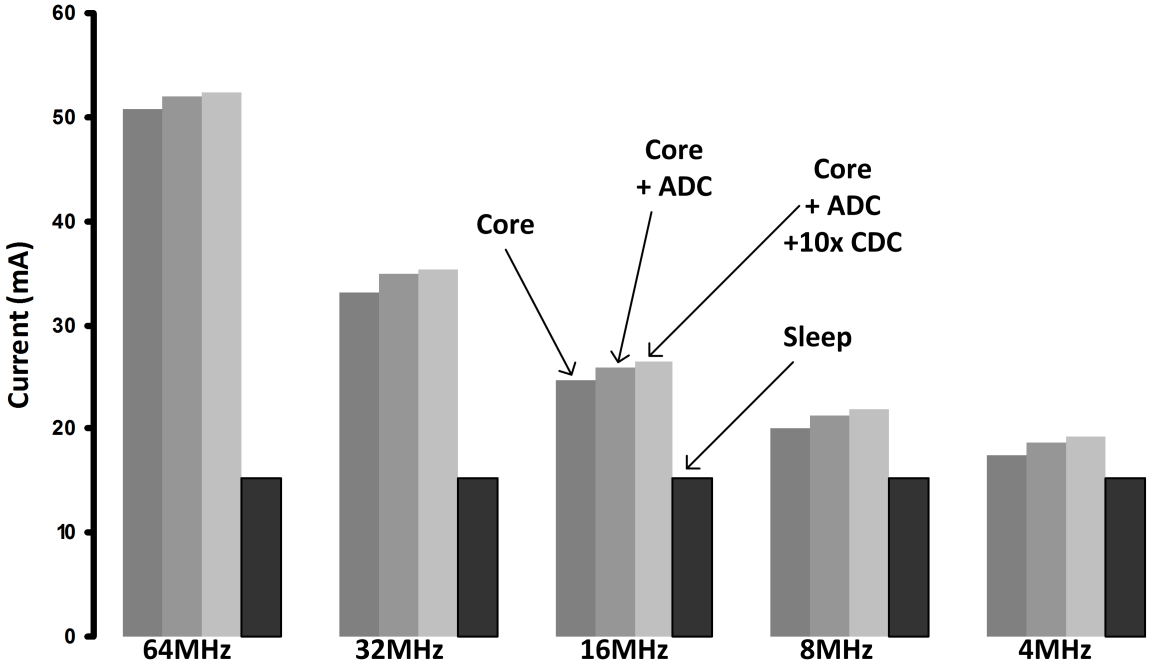

Figure 4.7: LIT's silicon measured current draw for active mode with operational blocks and Sleep mode at various frequencies.

In this section, we see how LIT has numerous strategies of lowering its energy consumption. 1) Using the Carbon-Zinc battery direct connect to allow it to use as much of the battery as possible. 2.) Dynamic Power Management Strategies used in LIT's power operational modes and Wakeup Interrupt Controller. 3) LIT's variable clock frequency and module power gating. These strategies allow LIT to achieve its low energy consumption lengthening its lifetime, achieving our low recurring cost goal. While these strategies are neither new to circuits nor system level implementations, the end result in energy consumption and its techniques to achieve it improves upon the current ICT solutions listed in 2.

4.4 LIT's Board Level Novel Configurable Pads

Unlike most general purpose processors or microcontrollers [24] [25] [26], LIT's LDOs are entirely on-chip thereby reducing the total number of external components in order to lower our cost. However, our choice of having on-chip LDOs limits our flexibility in voltage selection for our IO since we cannot configure IO pads to operate at particular voltage

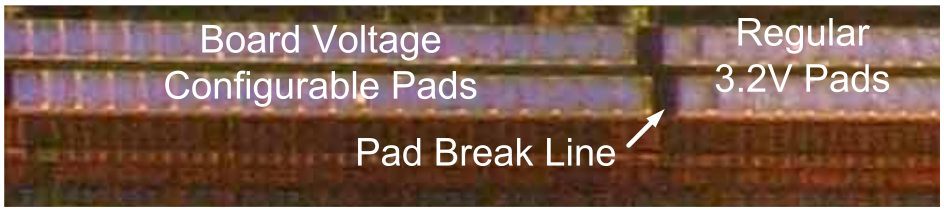

Figure 4.8: LIT's I/O Pad Break for board voltage configurable pads.

through software for communication with the off-chip components, such as a 1.8V or 3.2V NAND Flash, for example. Commercially microcontrollers would have a fixed IO pad voltage and would require an additional level converter chip in order to talk at the correct voltage levels. This means that additional chips would be required off-chip components, increasing LIT's final cost.

We overcome this by allowing the IO voltage to be configured at the board level (Figure 4.9). For example, if we have a 1.8V NAND Flash chip, we will connect the output of the Dirty LDO to LIT's IO pads' Dirty VDD. We accomplish this by having a break in the IO pads in LIT (Figure 4.8) in order to isolate the rest of the pads from the Board Configurable Pad's voltage domain. This is an extremely cost effective solution that requires only a single PCB trace to be wired from the output of either the Dirty LDO or Hybrid Switch Capacitor Network's output to the Dirty VDD of the Board Configurable Pads that is neither found in commercially available microcontrollers nor current ICT solutions.

4.5 Summary

In this chapter, we show how a very high level of integration of the power system of LIT results in a low off-chip component count that significantly reduces the cost to a point that is affordable by the targeted end users. LIT's ability bypass its LDOs and directly connect its Carbon-Zinc batteries result in a longer lifetime through eking out every as much charge as possible from the Carbon-Zinc batteries, power operational modes and Wakeup Interrupt Controller that minimizes energy consumption both meet LIT's goals of reducing its recurring costs. LIT further lowers its recurring cost through its variable clock frequency and module power gating by minimizing its energy consumption. LIT's Dirty LDO and Hybrid Switch

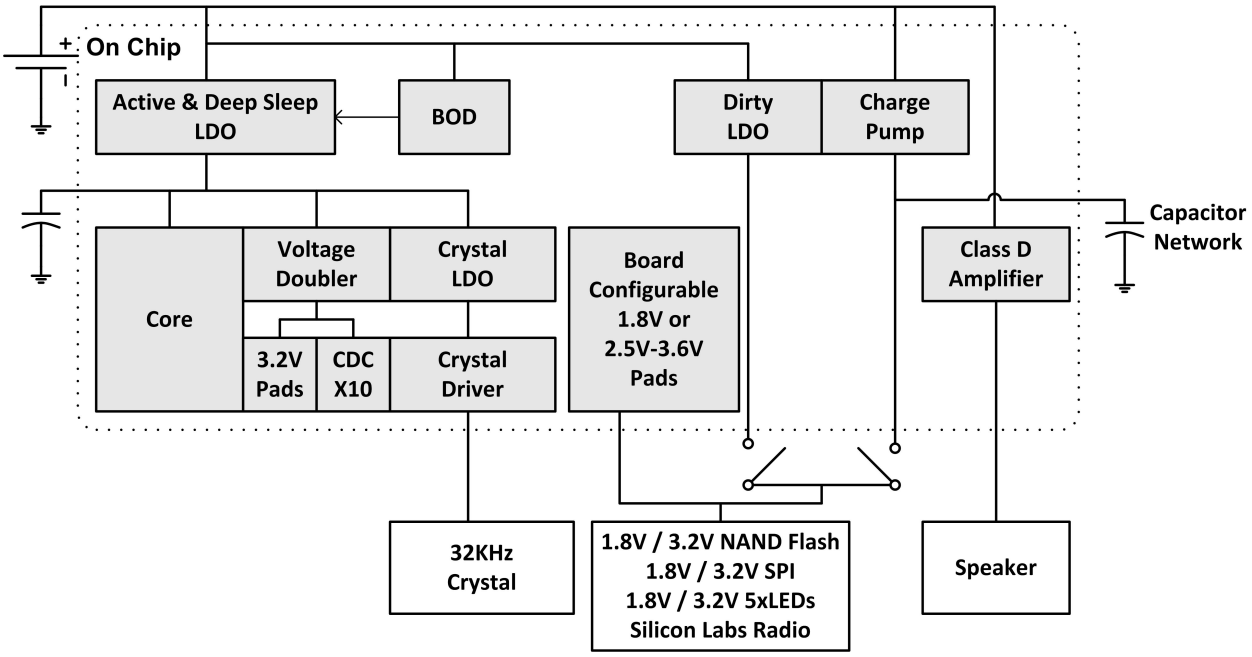

Figure 4.9: LIT's System Power Diagram: The NAND Flash Controller I/O can be powered by either the Variable Charge Pump or the Dirty LDO depending on whether a 3.2V or 1.8V NAND Flash chip is used. We accomplish this by connecting the Variable Charge Pump or Dirty LDO to the PAD Dirty VDD.

Capacitor Network in conjunction with its novel board level configurable pads allow LIT to be sensitive to the market volatility of its off-chip components, allowing us to choose between 1.8V or 3.2V components thus minimizing its initial cost. While the circuit techniques are not new, we showed how current commercially available solutions would require us to have multiple off-chip components in order to supply and regulate the power domains required by both LIT and its off-chip components. These commercially available solutions are already implemented in current ICT solutions and are a large reason as to why their cost is so prohibitive to these end users. These power strategies allow LIT to achieve its goal of a $6 initial cost while minimizing its recurring cost to the end users.

CHAPTER 5

Hybrid Switch Capacitor Network

5.1 Overview

The overall LIT system uses a small number of off-chip components that require a separate, constant supply voltage to operate properly (particularly, the NAND Flash chip and the Silicon Labs radio). Typically, these off-chip components require a supply voltage of 3.2V. However, the battery voltage has a large voltage range from 3.6 to 1.7V (two Carbon-Zinc batteries in series) and hence a regulation circuit is required to generate a constant 3.2V supply from the 3.6 - 1.7V battery voltage (Figure 5.1).

Typically, this is accomplished with an off-chip DC-DC regulation chip using a "boost converter" topology [54]. Such a boost converter requires a large inductor which is a relatively expensive discreet component (typical cost of $0.38). Also, the boost converter control system is complex which makes integration on the LIT chip more difficult. An alternate circuit using a switch-capacitor (SC) boost topology is simpler and uses much cheaper capacitances for its discreet components (typical cost of $0.04), making it more attractive for integration on the LIT chip. However, traditional SC boost circuits have very coarse regulation, making it difficult to meet the constant 3.2V output supply requirement given the highly varying battery input voltage. In this chapter, we introduce a new, hybrid "down-and-then-up" SC based topology that maintains the low cost and ease of on-chip integration of an SC boost converter but allows much finer voltage regulation, thereby making it possible to meet the constant 3.2V output voltage requirement. This novel SC topology is the result of the unique

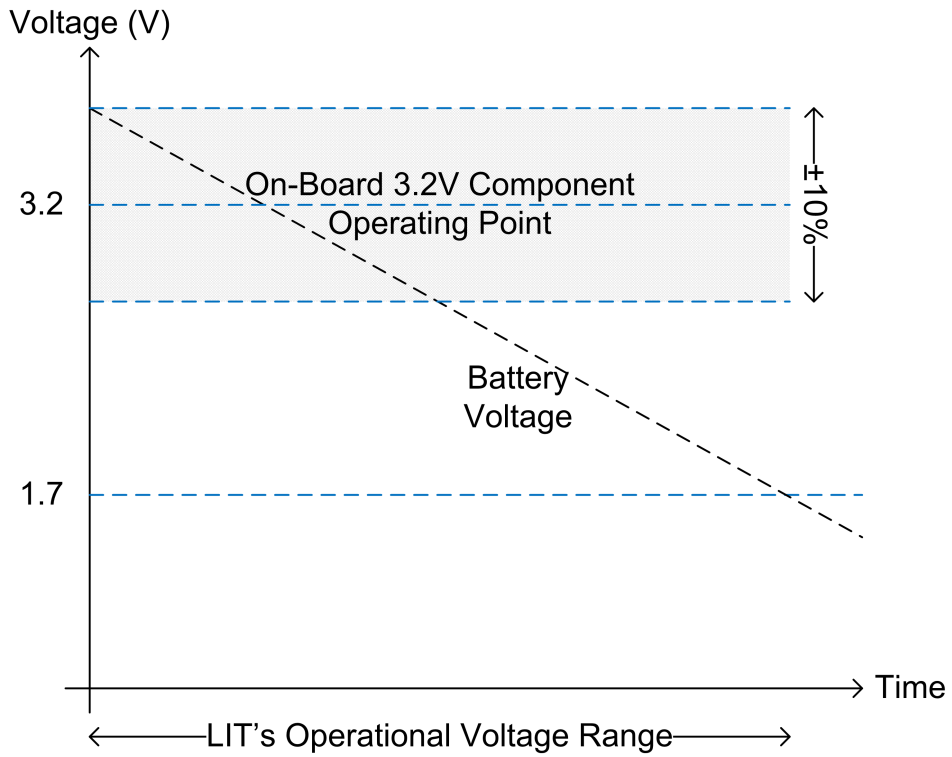

Figure 5.1: LIT's Off-Chip Component Operating Voltage Range.

requirements of LIT's particular application domain and demonstrates that new circuit level innovations can result from such requirements.

5.2 Prior Art in Voltage Converters

High efficiency voltage converters that are capable of both up and down regulation fall into two general classes: inductive converters and switch capacitor (SC) converters. We discuss each in turn and show how the unique requirements of LIT's application domain resulted in the need for a new converter topology.

Inductive converters can up-convert voltage using a boost converter topology (Figure 5.2) which typically consists of at least two switches and an inductor. The output voltage can be regulated through the pulse width and frequency that drives the switches [55] [56] [57] [58] [59] [60] [61] [62] [63] [64] [65]. Conversely, a buck converter topology can generate a voltage lower than its input voltage and similarly uses switches and an inductor [66] [67] [68] [69] [70] [71] [72] [73]. Again, the output voltage is regulated through pulse width and frequency

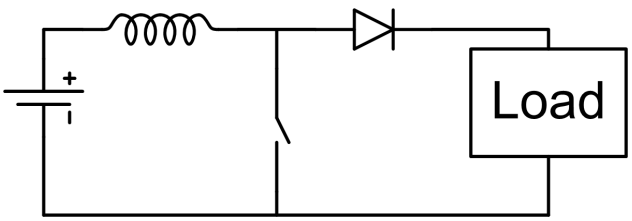

Figure 5.2: Typical Boost Converter Consisting of an Inductor, and Two Switches (Transistor and a Diode).

modulation. Simultaneously use of a buck converter followed by a boost converter will allow for a power converter that can supply a voltage output that is either lower or higher than the input. Regulation granularity can be very fine and is only limited by the pulse width control.

An inductive Buck/Boost DC-DC converter could have been used by LIT to satisfy the voltage regulation requirement for off-chip components. However, given LIT's low cost design goal through limiting the number and cost of off-chip components, inductors used by the Boost and Buck converters are disadvantageous. First, discreet inductors are more expensive than either resistors or capacitors, by at least one order of magnitude. Second, inductors tend to require more board area and height due to their size, further driving up the final cost of LIT when compared to capacitors. Third, we investigated by-passing these problems by implementing the inductor on-chip. However, on-chip inductors have low quality factors limiting the converter's efficiency. On-chip inductors also occupy a large area on silicon to achieve a reasonable inductive value, thereby increasing chip cost. Finally, boost/buck converters require fast and accurate pulse width and frequency control since they are very sensitive to load currents, which can change quickly. Hence, control must be implemented with fast, analog control circuits. Stability of these circuits is difficult, making on-chip integration a significant task and posing a practical obstacle for the LIT chip realization.

The second class of voltage regulations is Switched Capacitor (SC) DC-DC converters [74] [75] [76] [77] [78] [79] [80] [81] [82] [83]. These converters do not rely on inductors but instead use capacitors as their main energy storage element. Up conversion is typically obtained using a simple charge pump topology [84] [85] [86] [87] and relies only on capacitors. Hence, they are the cheaper choice when compared to Boost and Buck Converters. In addition, voltage control is accomplished through reconfiguration of the SC circuit (i.e, the number of SC

stages) and is not highly sensitive to load currents if the switching frequency is sufficiently high. Instead, reconfiguration is primarily used to respond to input voltage variation, which tends to be more slowly changing. This makes SC converters more amendable to simpler and slower control algorithms which can be implemented by the processor and do not require special fast analog circuits, reducing implementation complexity. However, SC charge pumps have an inherent problem of generating only coarse, fixed integer output voltage ratios of the input voltage of ∼1X, 2X, 3X, etc .... Clearly, such coarse conversion ratios are not sufficient to meet the 3.2V +/- 10% voltage regulation requirements of off-chip components.

To overcome this issue, we considered whether it was possible to overcome this difficulty by using a coarse SC charge pump followed by a fine grain resistive down converter (referred to as an LDO [36] [37] [38] [39] [40]). In this case, the charge pump would first up-convert the battery voltage to a voltage greater than 3.2V and the LDO would then down convert the voltage with high accuracy to 3.2V by simply dissipating the excess voltage. However, given the battery input voltage range, the LDO input voltage could be significantly higher than 3.2V. For instance, if the battery voltage was 3.0V, it would require at least doubling to ∼6.0V with the SC charge pump before the LDO would regulate it down to 3.2V. Similarly, with a battery voltage of 1.7V, the charge pump would likely need to triple the battery voltage to ∼5.1V since doubling the voltage to ∼3.4V would likely not allow for sufficient headroom for the LDO to achieve the required 3.2V. Such a high input voltage of the LDO has critical disadvantages which resulted in this approach being dismissed as a possible solution: 1) The voltage conversion efficiency would be significantly reduced (to $< 50\%$ based on simulation) since the LDO has a resistive operation. This in turn could increase operational cost of the device. 2) The high voltage would require a special high-voltage silicon process since the standard 180nm technology does not allow voltages higher than ∼3.6 volts. This would significantly increase the cost of the chip.

5.3 Proposed Novel Hybrid Switch Capacitor Network

In order to provide an output range of 3.2V over a wide range of battery input voltages from 3.6 - 1.7V, LIT proposes a novel SC topology that provides both fine grain voltage

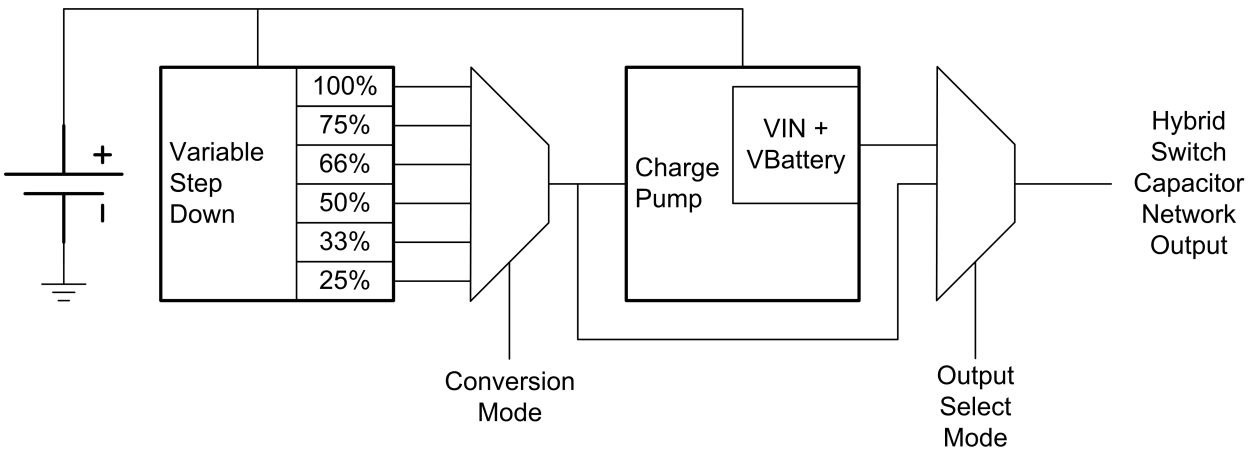

Figure 5.3: LIT's Hybrid Switch Capacitor Network Block Diagram.

regulation while maintaining a SC-only approach which has the advantage of using low cost capacitors and simpler control. The key idea is to first use a Step Down SC circuit, which can be designed to provide a variety of fine grain fractional voltages. In the LIT implementation, the SC down converter is designed to generate any one of the following fractional voltages: 25%, 33%, 50%, 66%, 75%, 100% of battery voltage(Table 5.1). This fractional voltage is then used to boost the battery voltage using a single stage, hybrid charge pump topology where the charge pump capacitors are driving by the fractional supply voltage instead of the battery input voltage. This hybrid topology results in an output voltage: Vout = Vbattery + Vfractional (Figure 5.3). Since the factional voltage has a fine grain resolution and high efficiency, the output voltage (Vout) can be tuned to within 10% of 3.2V over the input battery voltage range while maintaining high energy efficiency.

The Step Down SC converter achieves its fractional voltages of the Carbon-Zinc battery voltage input through configurations of its switches (Figure 5.4) (Table 5.1) which are chosen by the processor depending on what voltages the battery input is at, monitored by the on-chip ADC. Also, LIT's Hybrid Switch Capacitor Network can be bypassed to output the battery voltage directly or can only perform down conversion to power off-chip components with lower supply voltage requirements (e.g., 1.8V). The implementation is sensitive to energy consumption and can be power gated when unused. The Hybrid Switch Capacitor Network measures a relatively small 1.6mm X 0.7mm, which is largely due to its extremely large switches that has to be able to sustain up to 50mA current drawn from the external

components (Figure 5.6).

Table 5.1: Configurations for Step Down Switch Capacitor Network to achieve fractional battery voltages.

	25%	**33%**	**50%**	**66%**	**75%**	**100%**
ck1	Φ1	Φ1	Φ1	Φ1	Φ1	On
ck2	Φ2	Φ2	Φ2	Φ2	Φ2	On
ck3	Φ2	On	On	On	Φ1	On
ck4	Φ2	Φ2	On	Φ1	Φ1	On
ck5	Φ1	Φ1	Off	Φ2	Φ2	Off
ck6	Φ1	Φ1	Off	Φ2	Φ2	Off
ck7	Φ2	Φ2	On	Φ1	Φ1	On
ck8	Φ2	Φ2	On	Φ1	Φ1	On
ck9	Φ2	Φ2	Φ2	Φ2	Φ2	On
ck10	Φ1	Φ1	Φ1	Φ1	Φ1	Off
ck11	On	Off	On	Off	On	On

5.4 Detailed Operation

LIT's Hybrid Switch Capacitor Network operates in the following fashion, as explained by example. If the Carbon-Zinc Battery is at 2.0V, the proposed voltage converter would use a StepDown fractional voltage of 50% resulting in the output of the Step Down SC converter of ∼1.0V. Supplied into the Step Up SC converter, this would result in the Carbon-Zinc Battery Voltage pumped up to 3.0V. Figure 5.7 shows a simplified switch and capacitor topology for the 50% configuration where Φ1 and Φ2 are non-overlapping clocks generated on-chip from LIT's core clock frequency.

The fractional configuration for 50% output voltage can be seen in (Table 5.1). This results in the Step Down SC converter having the topology during Φ1 as seen in Figure 5.8 and during Φ2 as seen in Figure 5.9 which results in a Step Down SC converter output of 50% of the input voltage.

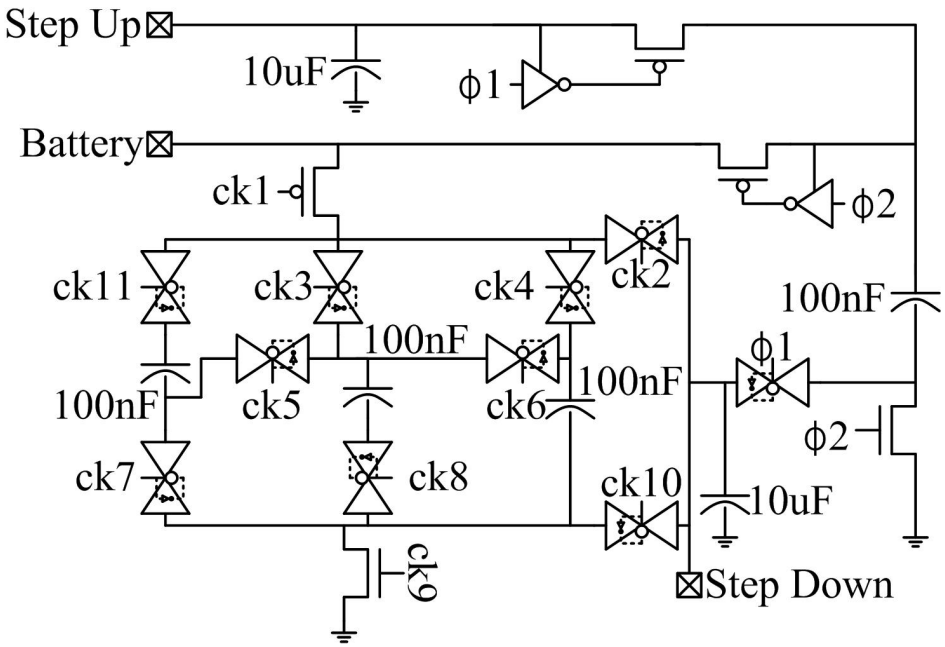

Figure 5.4: LIT's Hybrid Switch Capacitor Network.

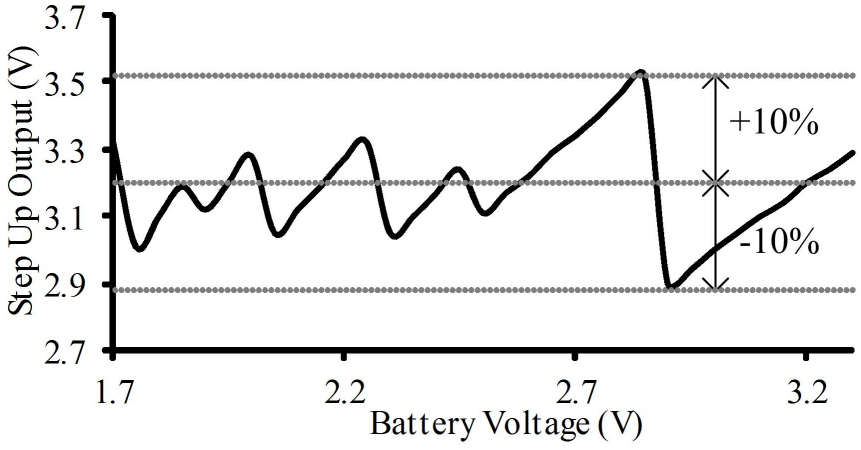

Figure 5.5: Measured Output of LIT's Hybrid Switch Capacitor Network is shown to stay within 10% of 3.2V across 3.3V - 1.7V battery voltage range.

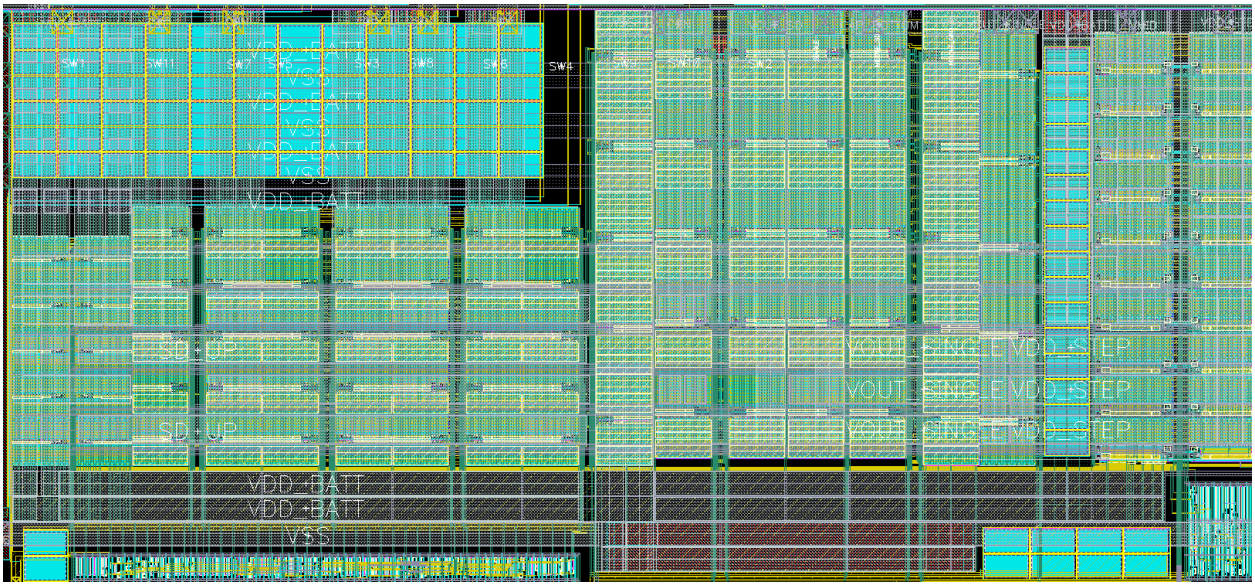

Figure 5.6: 1.6mm X 0.7mm Hybrid Switch Capacitor Network Layout.

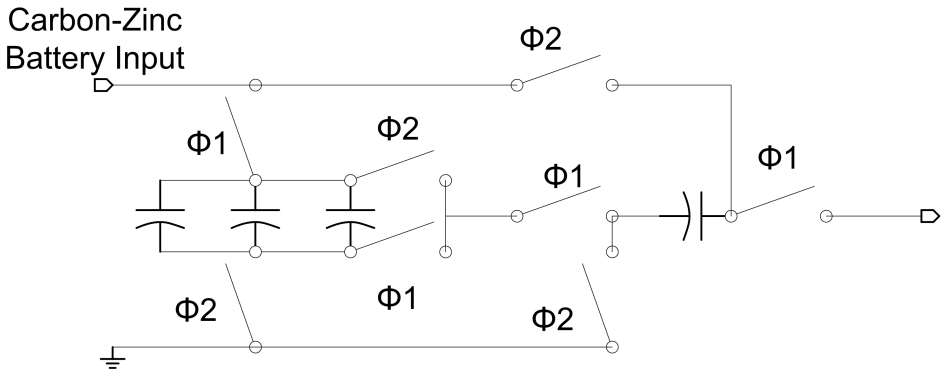

Figure 5.7: LIT's Simplified Hybrid Switch Capacitor Network Topology for 50% Configuration.

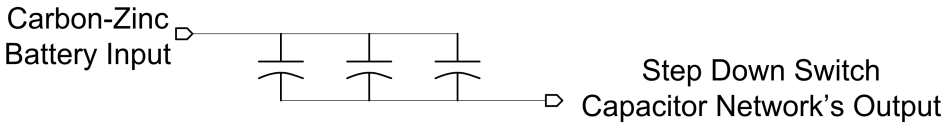

Figure 5.8: LIT's Step Down Switch Capacitor Network Topology during Φ1 for 50% Configuration.

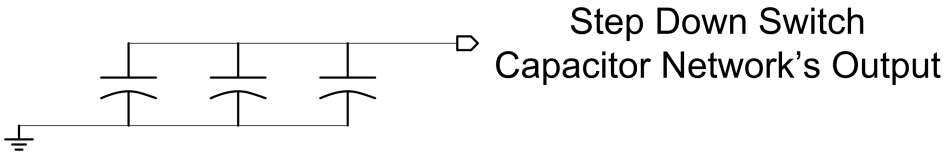

Figure 5.9: LIT's Step Down Switch Capacitor Network Topology during Φ2 for 50% Configuration.

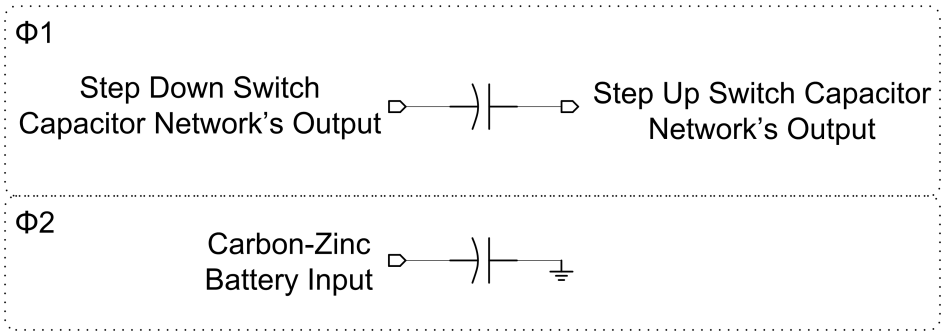

Figure 5.10: LIT's Step Up Switch Capacitor Network Topology for both Φ1 & Φ2.

5.5 Results

Figure 5.11 shows LIT's Hybrid Switch Capacitor Network's load regulation for various battery voltages, and Figure 5.12 shows LIT's Hybrid Switch Capacitor Network's line regulation. We see that the regulator with fractional Step Down and Step Up allows LIT to produce the required 3.2V for off-chip components.

5.6 Summary

In summary, LIT's novel Hybrid Switch Capacitor Network addresses discreet off-chip component cost by using only capacitors with a cost of $0.04, compared to the $0.38 cost of inductors used in boost/buck converters. Using a novel topology, it overcomes previous limitations of a switch capacitor converter's fixed conversion ratios by allowing the input of a SC charge pump to be driving by the fractional voltages of a SC Step Down converter. Silicon measurements show that the proposed converter is able to achieve regulation with 10% of 3.2V with up to 50mA of load current.

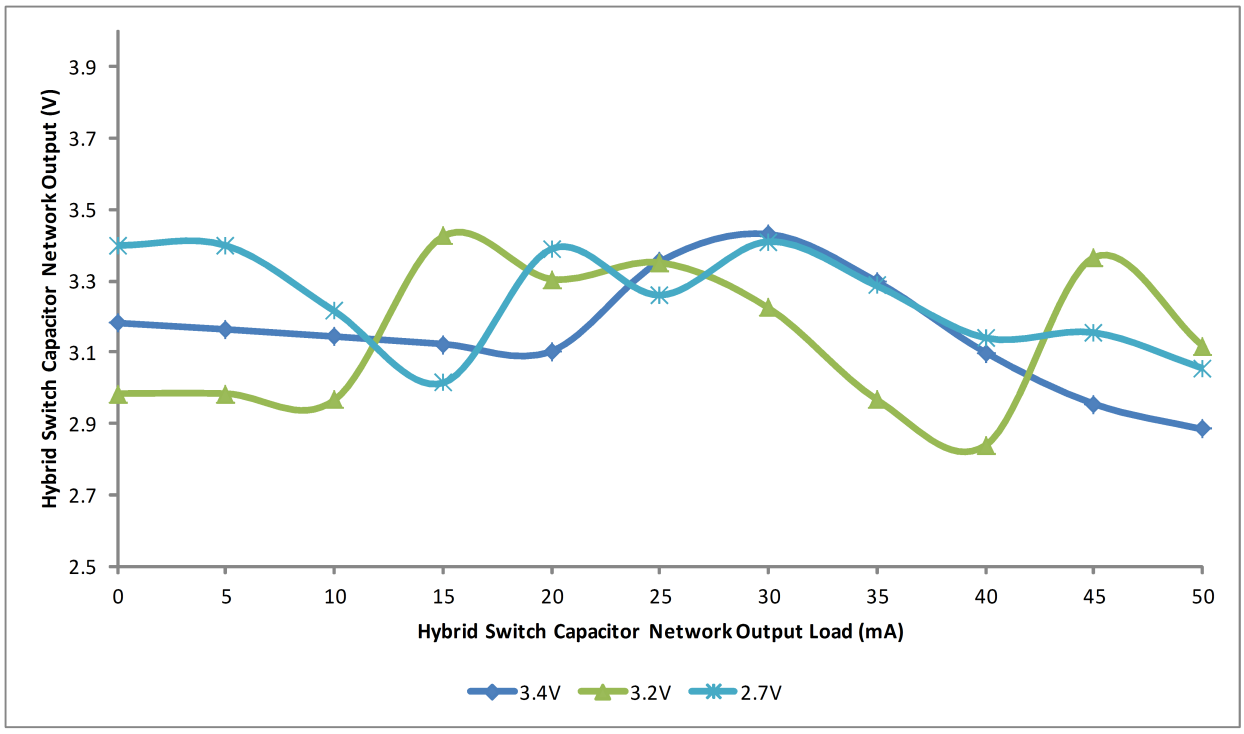

Figure 5.11: LIT's Hybrid Switch Capacitor Network's Load Regulation.

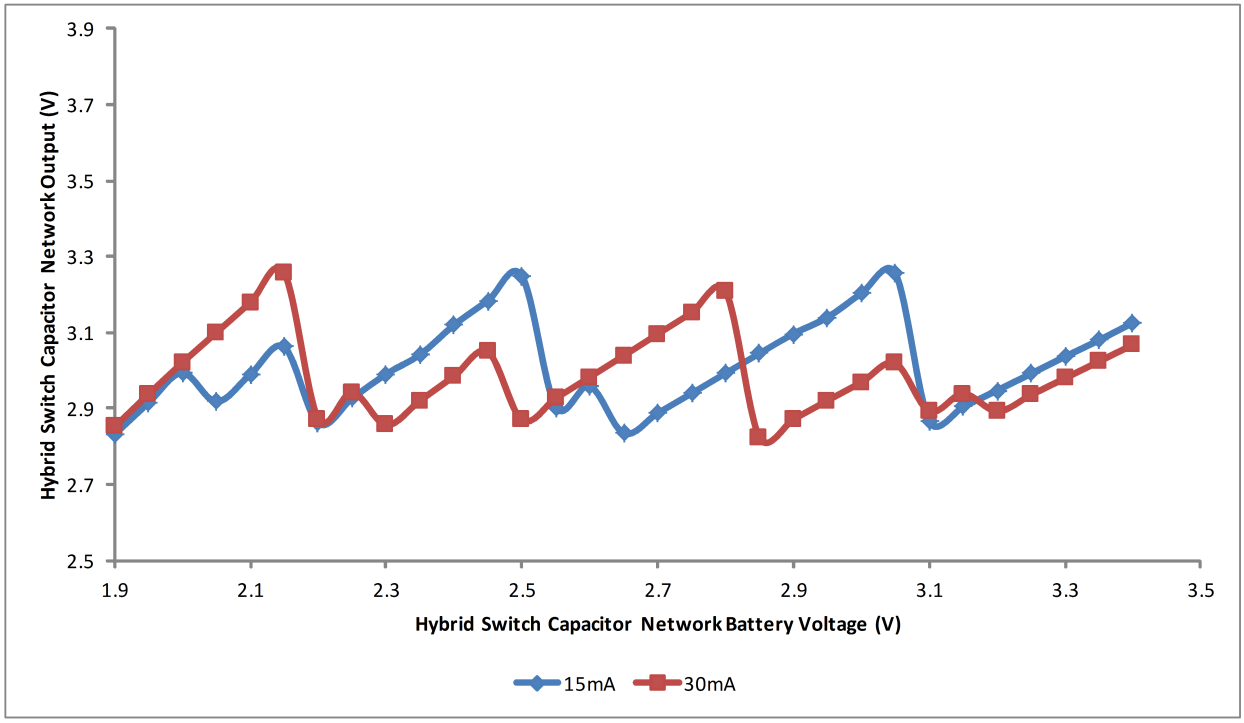

Figure 5.12: LIT's Hybrid Switch Capacitor Network's Line Regulation.

CHAPTER 6

Memory Hierarchy

6.1 Motivation

LIT has the unique problem constraint of achieving a low initial cost of $6 for total electronic as well as low recurring cost. We discussed various ways to achieve this in the thesis overview (Chapter 3), Integrated Power Management Systems (Chapter 4), and Hybrid Switch Capacitor Network (Chapter 5). While LIT's memory hierarchy also has this same cost constraints to satisfy, it also has the constraint of having to provide a large memory space in order to store the audio data that is needed alongside with its code base. In order to achieve this, we developed a 128kB 4-way True LRU design backed by a 2Gb off-chip NAND Flash chip that allows LIT to meet its initial cost target and minimize its recurring cost all while maintaining a large memory space with reasonable performance.

In this chapter, we discuss the commercially available solutions and why they do not fit our constraints. We will then present other solutions investigated, their shortcomings, and show how we overcome LIT's unique constraints with the proposed memory hierarchy. However, the consequences of using this memory hierarchy in this system results in several unusual obstacles that need to be addressed. We therefore also present how we implemented this memory hierarchy and its results and finally present our concluding remarks.

6.2 Commercially Available Microcontrollers, Memory Hierarchies and their Cost

We initially began our search for an ICT solution by looking extensively at currently available microcontrollers that could potentially fit the need of our application while meeting our goal of a low cost system with sufficient memory size and performance. We found that commercially available solutions were either too small from a memory space standpoint, or well beyond the price range our application can bear. Once we came to the conclusion that current microcontrollers did not present a feasible solution, we also looked at other possible memory solutions that would yield a cheap solution which could be integrated with LIT.

Currently commercially available smaller microcontrollers such as TI's MSP430G2001 [88] ($0.39), other ARM M0 based microcontrollers such as NXP's LPC1110 [89] ($0.823), and Silicon Labs' C8051T600 [90] ($0.675) are within the price range for the application space that LIT is in. However, their memory space is very small. As shown in Table 6.1, TI's MSP430G2001 only has 0.125kB of RAM with 0.5kB of Flash, NXP's LPC1110 only has 4kB of RAM with 16kB of Flash, Silicon Lab's C8051T600 only has 256Bytes of RAM with 8kB of Flash. Their small memory space is a significant limitation for our application space requiring additional off-chip memory components in order to hold the required data which eventually would add to the cost of the system.

Table 6.1: Currently available microcontrollers, their memory space, and cost.

Microcontroller	RAM	Flash	Cost
TI MSP430G2001	0.125KB	0.5KB	$0.39
NXP LPC1110	4KB	16KB	$0.823
Silicon Lab C8051T600	0.256KB	8KB	$0.675
TI RM42L432	32KB	384KB	$6.53
Atmel ATSAM4S16B	128KB	1MB	$5.15

Higher performance microcontrollers such as TI's RM42L432 [91] have a larger memory space with 32kB of RAM, 384kB of Program Flash which makes it more desirable from a memory size standpoint, but has a cost of $6.53, which by itself exceeds the total electronic cost goal of $6. Other similar choices such as Atmel's ATSAM4S16B [92] has 128kB of RAM with 1MB of Flash, but, similarly, has a prohibitive cost of $5.15. The reason for the high

cost of these higher-end micro controllers is that they use a larger and higher performance microprocessors (such as ARM Cortex-R4 and M4 class architectures) than the low end microcontrollers (which use ARM Cortex-M0 class architectures). This adds significant cost for additional functionality and performance which is not strictly necessary in our application domain. In addition, these microcontrollers have significant on-chip embedded flash arrays. Since embedded flash ($1F^2$ —$100F^2$) [93] is less dense than standalone flash (NAND: $4F^2$, NOR: $10F^2$) [94] [95], it is more economical to place all flash storage on a separate flash chip. Furthermore, while these microcontrollers have larger SRAM and Program Flashes, they still would require additional memory space to store the amount of audio data that the system requires.

These high performance microcontrollers also typically have a NOR Flash for program data and DRAM for volatile memory since NOR Flash is slow and on-chip SRAM is still limited. NOR Flash is typically used over NAND Flash for program storage due to its ability to do fast random memory accesses which makes it easy to use from a hardware perspective since it can be hooked up similar to memory and accessed directly. However, NOR Flash is typically less than 50% as dense as NAND Flash which results in being more expensive from a cost/bit perspective. NOR Flash is also more power hungry, consuming ∼15 mA than NAND Flash chips ∼10mA [96] [97] [94]. Because of its high-speed, direct access capability, NOR Flash requires separate pins for every address and data bit which can result in as much as 26 additional pins on the LIT package, which increases the cost of packaging. This is unlike NAND Flash which typically uses an 8-bit or 16-bit bus and overloads their functionality for both commands and data input/output.

Given the limitations of currently available microcontrollers and their memory hierarchies, we opted to build LITs memory hierarchy from the ground up, removing the expensive components typically found in the current ICT solutions or typical microcontroller's memory hierarchies. From our investigations of Talking Book's Bill of Materials, we found that the NOR Flash used to store the codebase (∼$0.84), and the MicroSD (∼$3.00) used to store the data were the two most expensive components of the overall system excluding the general purpose processor and PCB. This prompted us to investigate other possible solutions for the memory storage needed for code and data. This included investigation of other possible

sources of cheaper MicroSD chips, on-chip embedded flash, and a single unified NAND Flash solution.

We initially began our investigation with recycled MicroSD solutions from ReCellular [98], the world's largest secondary wireless industry, that collects and recycles unused cell-phones. We attempted to obtain recycled MicroSD cards through ReCellular through their donated phones, but this proved to be an unfeasible solution because 1) The MicroSDs received were not all the same type or size. 2) The MicroSDs received were not always in good condition and had robustness issues after being in use for so long. 3) The cost of used MicroSDs was still considerable (∼$2.00) since there is an active market for them overseas. This resulted in us abandoning the recycled MicroSDs as a viable option.

The second possible solution was to use on-chip embedded flash. This was investigated but was found to be too costly in terms of die area, since embedded flash is ∼4X less dense then standalone flash. Also, the technology we selected for LIT, TSMC180, is not an advanced technology node and therefore further increases the area and cost of a large embedded memory. While a smaller technology node might have yielded better results, the goal of the integration of all the analog components onto a single chip would have been much more complex due to the lower voltage headroom in the newer technologies. Also, newer technologies have much higher initial cost for mask generation. This prompted us to look towards off-chip solutions. While an off-chip solution is counter to LIT's design goals of a high level of integration, the memory hierarchy is one of the few areas where having components off-chip results in a lower cost than if we had implemented it on-chip.

Based on our previous findings, we concluded that we could only meet the constraints by combining the following components: a) Large off-chip NAND flash chip (low cost and low standby power), b) a low power, low cost small micro controller, such as M0 (other higher end processors are large and expensive and more power hungry). c) a large on-chip cache to shield processor from slow NAND flash and improve performance / power. Note how this meets our overarching goal of a low number of off-chip components in order to reduce cost. Furthermore NAND Flash is relatively cheap in terms of cost (∼$1—-$2)and power (10mA), compared to NOR Flash, MicroSD, or Embedded Flash. This results in a reasonable chip area for the processor and cache with a low pin count, allows for a performance of 10 to 100

MHz, and compact code size.

However, there are several consequences / obstacles that arise when creating a working system with the somewhat unusual combination of these three components, such as requiring an FTL running on the chip, software cache miss and fill handling using interrupts, loading FTL from the NAND Flash, a large cache size for the code space, and a configurable pinned region that led to a True LRU cache. All of these, we explain in the following section.

6.3 Memory Hierarchy

LIT's memory hierarchy is designed to accommodate the need for large code space (up to 16MB) and data space while eliminating the need for expensive and power hungry NOR Flash, DRAM, and MicroSD commonly in microcontrollers and current ICT solutions. As mentioned, LIT implemented a large 128kB 4-way True LRU on-chip cache (48% of die area) directly backed by an off-chip NAND Flash Chip (Figure 6.1).

One of the complexities of implementing a NAND Flash chip is that it raises the need to implement a Flash Translation Layer (FTL) to interface with the flash and perform load leveling and hide the block-wise erase and write operations. However, since this FTL has a high amount of complexity, the cache fill can no longer be performed in hardware and must be performed in software. Hence, we have a hardware enabled miss detection with a fill mechanism that is handled through software upon a miss. Since the processor of choice, the Cortex M0, does not have a proper memory management unit such as a page system, LIT performs the software cache fill using interrupts that was not their original intended purpose.

Since the FTL is too large and too prone to changes and updates to reside in a permanent and unchangeable ROM, we developed a boot system where we initially boot from the NAND Flash without the FTL in order to load the FTL itself. Furthermore, NAND Flash speeds are slower than DRAM speeds and because of that, we chose to have a large 128kB of on-chip cache implemented with SRAM, which consumed 48% of total die area. Although consuming almost half the total die area, this only added ∼$0.25 to the total die cost, which is still significantly cheaper than having DRAM and guaranteed enough code space for our

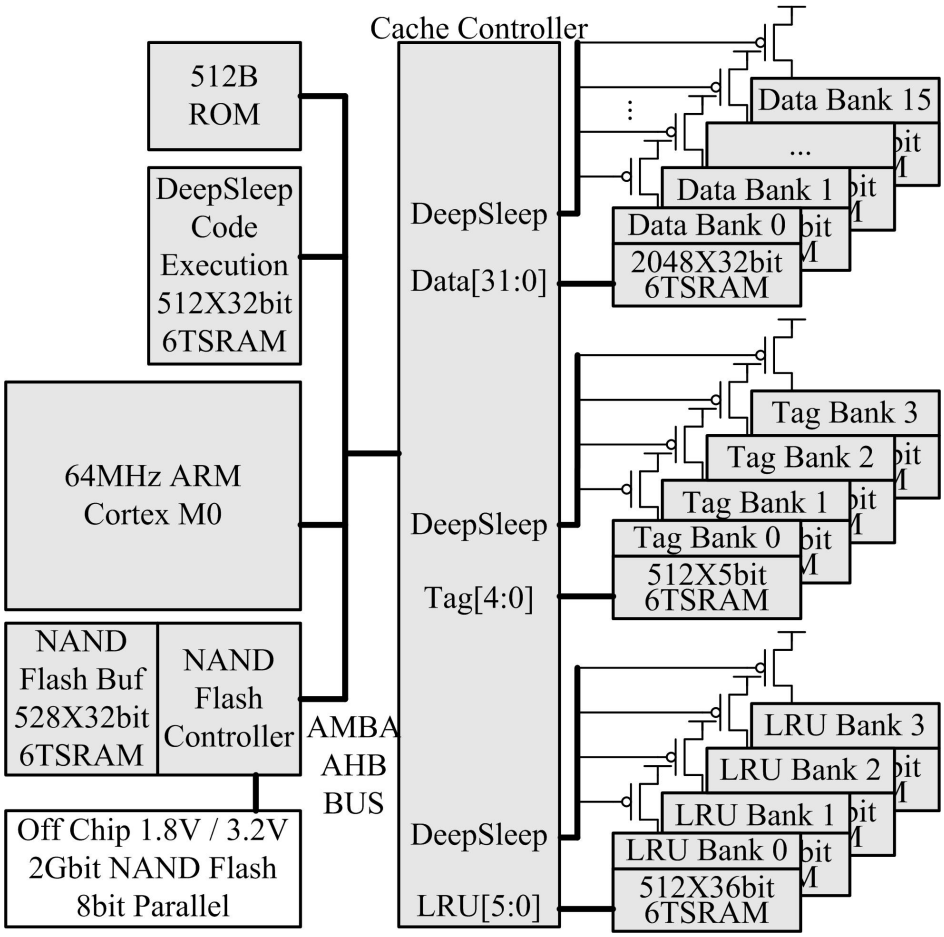

Figure 6.1: LIT's memory system diagram.

application.

The cache has a configurable pinned section that prevents a lower portion of the address space from being evicted. This is necessary for storing the OS and FTL code since we need to guarantee that fault handling code is never evicted, else we would not have any way to fill the cache upon a miss. The pinned section is achieved by having each cache access checked against a configurable address pin line in hardware to determine where it resides. Upon a miss, the cache causes a precise fault which alerts the core and executes code in the fault handler that then loads data from the NAND Flash into the cache.

A further consequence of the pinned region is that we must use a True LRU cache line replacement policy in hardware since pseudo LRU cannot provide a fallback address to evict if the returned LRU-way is in the pinned region. If the returned LRU-way is in the pinned region, we need to know what the second least recently used address is in order to evict it. Pseudo LRU only provides the least recently used address, raising the need for True LRU. If the second least recently used address is also in the pinned region, we need to know what the third least recently used address is which True LRU can provides as well.

Another consequence of backing the cache with NAND Flash is that we cannot simply write back upon eviction from the cache, since programs that cause cache thrashing would wear out the NAND Flash very quickly. To address this, the cache uses an explicit data write back architecture where we only write back data when we absolutely need to retain it. Hence, modifiable data is also stored in the pinned section. This way, if important data needs to be evicted, software knows when it is evicted and whether it needs to be written back or whether it can be discarded if it is unneeded.

We see that our goal of designing a low cost system, which resulted in a large off-chip NAND flash chip with a small micro controller, and a large on-chip cache, has resulted in additional complexities in the system. These include the FTL, hardware miss detection and interrupt driven software miss filling by the CortexM0, a pinned region to guarantee fault handling code, True LRU for fallback addresses to evict, and cache thrashing prevention. These solutions overcome the system complexities which allow us to achieve our low cost and robustness goals.

6.4 Memory Operation and Uniqueness of LIT's Memory Hierarchy

Upon powering up, a 512B boot ROM on LIT automatically fills the cache with the software miss handler and FTL in the pinned section ensuring that they will be available for future operations. LIT takes 1s to come out of boot and perform the required operations to put LIT at a point where the user can start interfacing with it. The long boot time is due to our making LIT more robust from a process variation standpoint. LIT initializes all its timing configurations to the slowest possible points to account for process variation, and LIT is then sped up through configurations stored in the NAND Flash. The configurations in the NAND Flash are also used to initialize peripherals. By leaving configurations on the NAND Flash, LIT is more flexible and robust. After booting, code is then executed out of the cache and misses are handled as previously described.

After some idle time when software has determined that the device is not in use, LIT enters Deep Sleep mode as described in Section 4.3.2. In order to exit Deep Sleep mode and re-enter Active mode, LIT stores a piece of code, the event validity checker, in memory that is executed occasionally. The event validity checker serves two purposes. First, it is used to validate the requests from the capacitive sensors that we use as the human input interface since false positives could wake the entire core unnecessarily, waste energy, and produce unwanted behavior reducing its robustness. Second, it is used to track changes of the capacitive sensors and reconfigure its thresholds as temperature, background noise, and variation will result in its value drifting slowly over time, which could, again, result in unwanted behavior.

In Deep Sleep mode, LIT has two options for maintaining the event validity checker in memory. The event validity checker can be stored in a dedicated 2kB of SRAM (DeepSleep Code Execution Figure 6.1). In this case, LIT power gates the entire cache and saves the most energy. Alternatively, we can store the event validity checker in a section of the cache that is purposely left awake. This allows more complicated event checking code that is unable to fit in the 2kB of SRAM. In this option, we are capable of maintaining a larger, variable amount of active memory during Deep Sleep mode with a granularity of 8kB. Upon

the event validity checker verifying an event, the cache refill routine that is already in the boot ROM will be executed and LIT will re-enter Active mode. If the event validity checker determines that the event is a false positive and therefore invalid, LIT will simply goes back to Deep Sleep mode without waking the large cache (Figure 6.2). The event validity checker, combined with the granular power allows LIT to minimize its energy consumption in Deep Sleep mode.

The techniques used in the memory hierarchy are not new in-and-of themselves. For instance, the use of hardware miss detection is common in most high-end processors and the use of software miss handling has also been proposed [99] [100]. Park [99], presented a cost-efficient memory architecture backed by an off-chip NAND Flash chip. However, their system boasts a ARM A9 core with an L1 cache and full memory management capabilities which allows the miss filling to be performed in hardware. LIT's simpler CortexM0 requires a software miss fill capability since it does not have memory management capabilities. Furthermore, there is no discussion about the complexities and overhead of implementing an FTL on their system, which LIT does. From a system level perspective, the particular combination used in the LIT hierarchy (a large off-chip NAND flash, large on-chip SRAM and a small, low power processor) is unique in the market place (Table 6.1). To our knowledge, LIT is the first to use this particular combination of components, which we believe is critical in obtaining the required cost savings. Particularly in LITs application space, the system is novel compared to the ICT devices listed in Chapter 2, which all use some combination of MicroSD cards, NOR flash, and DRAM.

6.5 Memory Hierarchy's Energy Consumption Results

In total, LIT's 128kB cache draws an additional, silicon measured, 1μA in Deep Sleep if one 8kB of memory is left awake and 7.2μA if all 128kB is left awake (Table 6.2). The disproportionally higher current draw for having the first 8kB left awake is due to the need for all LRU and Tag banks to be left active during Deep Sleep mode to ensure that we know what data is valid. Each additional 8kB of active memory results in an additional 0.4μA in current draw. This indicates that the LRU and Tag banks draw a total of 0.6μA. Having the

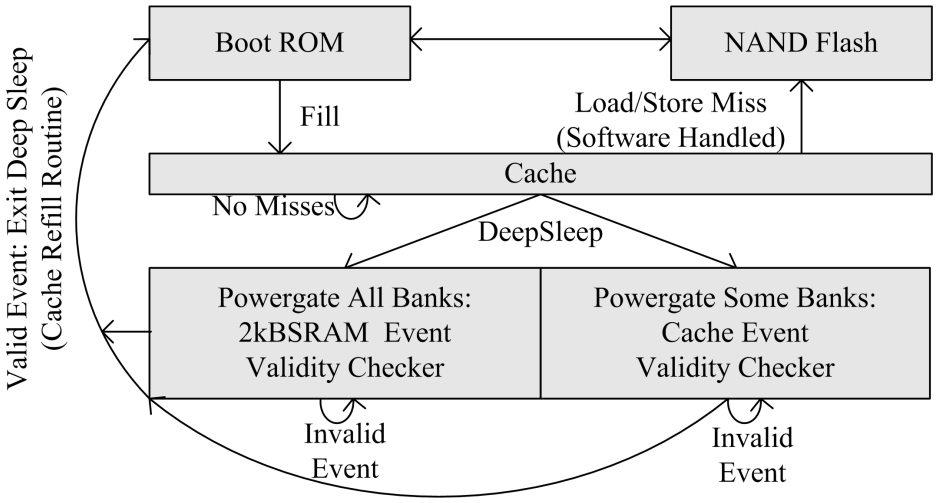

Figure 6.2: LIT's memory operation from loading from boot, miss handling during Active mode and Deep Sleep mode cache power gating options or using 2kB dedicated SRAM.

ability to control the amount of memory left awake during Deep Sleep mode allows software to be more flexible in terms of the complexity of their event validity checker. Software has control over just how much code is needed and can adjust the number of data banks to be power gated in order to consume as little energy as possible in Deep Sleep mode. When all 128kB of memory is disabled and only the 2k event SRAM is left on, the standby power is reduced to only 17uA.

Table 6.2: Silicon measured current of LIT's cache during Deep Sleep mode.

Deep Sleep Cache Memory Awake	**Additional Current**	**Total Deep Sleep Current**
0kB	0μA	17.0μA
8kB	1.0μA	18.0μA
16kB	1.4μA	18.4μA
32kB	2.2μA	19.2μA
64kB	3.9μA	20.9μA
128kB	7.2μA	24.2μA

6.6 Summary

In conclusion, we show how current embedded microcontrollers and ICT solutions have memory hierarchies that cannot meet LIT's design goals of a low initial cost, large mem-

ory space and reasonable performance. Beginning with the investigations of other possible memory solutions, we found that by using a large off-chip NAND Flash with a large on-chip 128kB 4-way True LRU cache in conjunction with a small, low power low cost microprocessor (cortex M0), LIT manages to satisfy the cost and performance constraints through reducing the number of off-chip components and the minimization of energy consumption targeted towards our application. Despite the additional obstacles and complexities that were borne out of this memory hierarchy, we manage to overcome these complexities. We achieved a low energy consumption through having an event validity checker that allows the entire cache (or parts of the cache) to be power gated during Deep Sleep mode. This allows flexibility in the amount of code that it needs to be kept available. The event validity checker also increases the power efficiency of the device by not requiring LIT to completely move into Active mode during false positives of the capacitive touch buttons. This reduces the recurring cost to the end user. These design decisions and solutions show how we can build a low cost system not achieved by the previous ICT solutions listed, with a low cost, low power memory hierarchy targeted towards our application space.

CHAPTER 7

Power-on-Reset / Brown-out-Detector with Lock-off for Carbon-Zinc Batteries

7.1 Introduction

In Chapter 3 we introduced LIT which is a custom silicon device designed for operation with a single carbon zinc battery input. In this chapter, we discuss in more detail the Power-on-Reset (POR) / Brown-out-Detection (BOD) circuit which was specifically designed for the LIT chip, operating on Carbon-Zinc batteries. Since Carbon-Zinc batteries have a strong "healing" behavior, a much larger hysteresis is required than a traditional POR/BOD circuit can provide. Hence, we present a new circuit which introduces a so-called "lock-off" function which allows hysteresis to be increased dramatically and prevents the oscillations that we observe with a traditional circuit. Below, we provide an overview of the POR/BOD operation, review prior work and propose our new method which is uniquely designed for the LIT context.

7.2 Background and Unique Requirements for LIT Context

A key requirement for any battery operated system is that the system functions predictably when the battery is first inserted and when the battery is discharged and its voltage

drops below a safe operable level. To address this, electronic systems are equipped with a (POR) / (BOD) circuit. The POR ensures that the system is properly reset after the battery is first inserted. When power is initially applied, the values stored in the sequential logic gates (FlipFlops, latches, and memory) are in an unknown state, which can lead to unpredictable behavior. To circumvent this, the reset signal is asserted by the POR circuit while the battery voltage ramps to a safe level, guaranteeing a known initial state from which the device commences operation. Similarly, when the battery is discharged to the point where the voltage is too low for reliable operation, reset must again be asserted by the BOD circuit to prevent unpredictable behavior at the low operating voltage and could possibly corrupt its state. Typically, the function of POR and BOD are combined in a single circuit.

LIT's application space has the unique problem of using batteries from developing regions, typically consisting of Carbon-Zinc batteries. Even though Carbon-Zinc batteries hold a much lower charge than the alkaline batteries ($\sim$600mAh vs. $\sim$2000mAh) they are prevalent in developing regions for two reasons: 1) Their purchase cost is significantly lower than that alkaline batteries, making them more accessible to low income populations, even though their cost per mAh is possibly higher [1]. 2) They have a very long shelf life which is critical in remote, low income regions where transportation of products can be very slow.

Carbon-Zinc batteries present the challenge in that the voltage of the battery exhibits a very strong self-healing property when unused due to its particular chemistry. During our testing of Carbon-Zinc batteries, we found that when current is being drawn, the voltage of the battery will slowly degrade, as expected. However, when current draw stops, the battery's voltage starts to heal, causing the battery voltage to increase significantly. Figure 7.1 shows how the voltage decreases and increase when the power draw is pulsed, showing as much as $\sim$700mV of self-healing.

This large swing of battery voltage poses a unique problem from a reset perspective since traditional POR/BOD systems have a hysteresis which is on the order of $\sim$100mV, which is smaller than the self-healing voltage swing. If during system operation, the voltage of the Carbon-Zinc battery drops below its minimum operating point, a traditional POR/BOD

[1] note that since alkaline batteries were not available in remote, low income locations, exact price comparisons were difficult to quantify

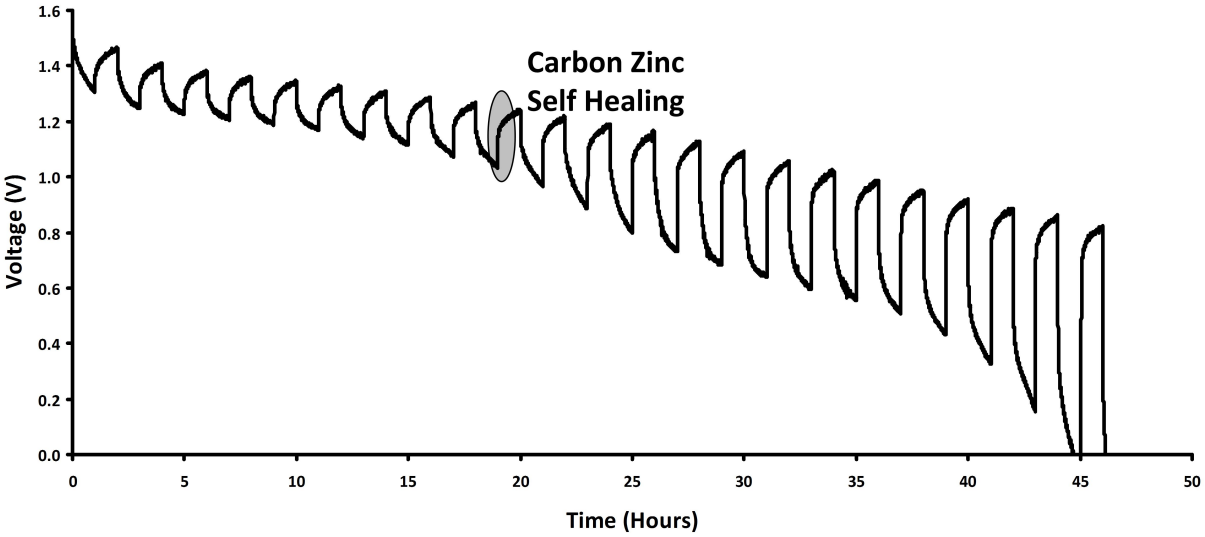

Figure 7.1: Ghana's Carbon-Zinc batteries exhibiting its self healing property.

circuit will assert reset to stop the system from operating further. However, this will cause the voltage of the battery to increase due to significantly reduced current and the self-healing properties of this battery chemistry. The battery voltage will then reach a voltage where the Power-on-Reset / Brown-out-Detector de-asserts reset and enables the system thus consuming more current thereby driving the voltage low again.

(Figure 7.1) shows how this increase and decrease in current draw leads to an oscillation that will continuously turn the device on and off near the end life of the battery. The voltage waveform is the actual measured behavior of a talking book device which was utilizing a traditional POR/BOD circuit. Clearly, the oscillating behavior is highly unwanted since it: a) presents confusing behavior to the user at the end of battery life as it creates a buzzing sound which leads users to conclude that the device is broken and b) increases the likelihood of NAND Flash corruption resulting in loss of data due to repeated powering on and off of the NAND flash chip.

7.3 Prior Work in Power-on-Reset / Brown-out-Detectors

Figure 7.2 shows a simple traditional Power-on-Reset circuit consisting of a RC component connected to an output inverter driving a reset signal. Initially, when the battery

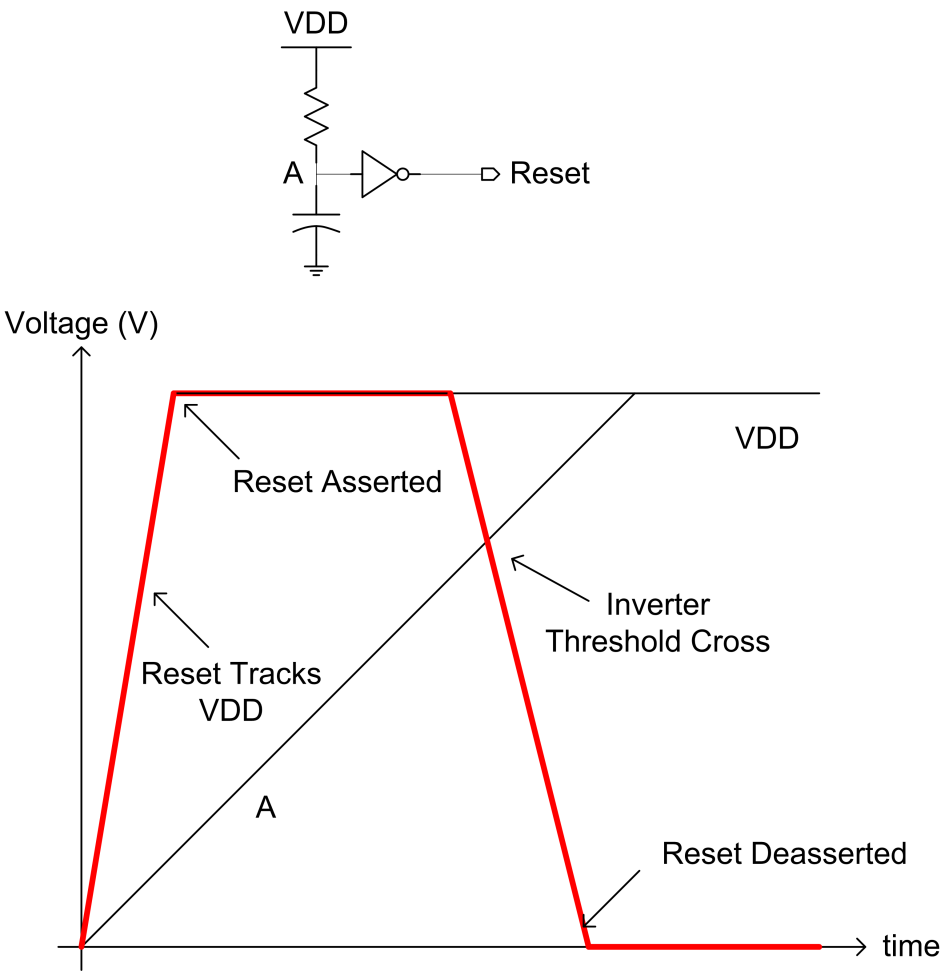

Figure 7.2: Traditional Power-on-Reset Circuit.

voltage (i.e., VDD of the system) is ramping up, the RC component ensures that node A will rise slower than VDD and hold the output reset driver at VDD until node A crosses the inverter's threshold. This ensures proper startup behavior when VDD is initially applied. Furthermore, it can be easily implemented with a few off-chip components. However, the RC time constant (on the order of milliseconds) makes integration on-chip difficult since it requires both a large capacitance and resistance, both of which require large silicon area (e.g., a 10ms RC constant could be implemented with a 10pF capacitance and 1Gohm resistance, which would require 1.2mm^2 in a 180um process technology which is unfeasible).

More sophisticated approaches for POR/BOD have focused on on-chip implementations of low power implementations for embedded portable device systems [101], robustness from a temperature and process variation aspect [102] [103], brown-out-detection with adjustable

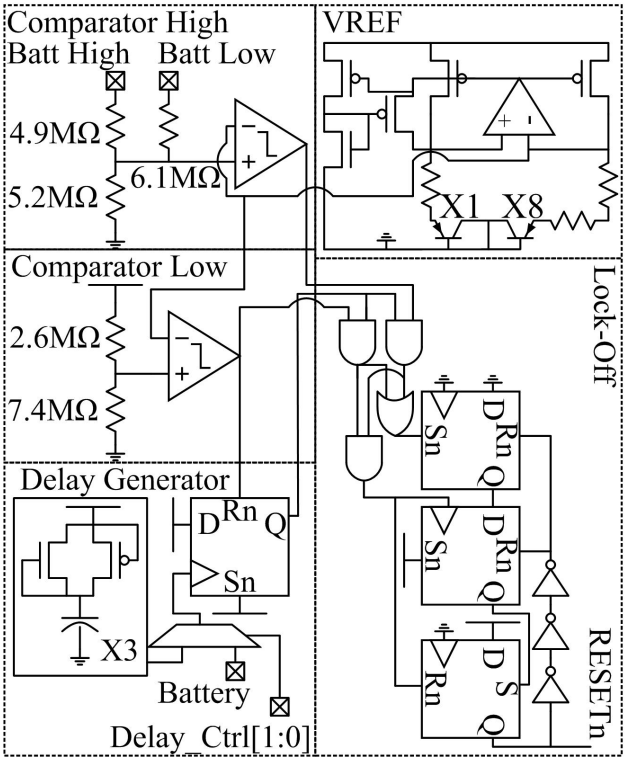

Figure 7.3: LIT's on-chip 1.7μA Power-on-Reset / Brown-out-Detector with "lock-off" for Carbon-Zinc Batteries.

thresholds of brown-out ramp time [104], adjustable reset times for slow voltage ramps [105] [106], and low power supplies with power supply immunity [107]. However, all previous on-chip approaches have a hysteresis that is on the order of 15mV - 144mV [108] [109] making them unsuitable for LIT's application space. This is due to the fact that the hysteresis of these approaches is derived from the innate behavior of an analog comparator. Since in traditional applications, a large hysteresis is considered undesirable, no research on on-chip POR / BOD circuits with a high hysteresis were found.

7.4 Proposed Novel On-Chip POR/BOD Circuit

We propose the following novel Power-on-Reset / Brown-out-Detector with a "lock-off" feature to increase the hysteresis such that the undesired this oscillation at end-of-life can be avoided with Carbon-Zinc batteries (Figure 7.3). The circuit operates as follows: Upon battery insertion (Figure 7.4), the supply voltage is low enough such that both the 1.7V

(Comparator High) and 1.2V (Comparator Low) Comparators' outputs are low, making the outputs of the 1st and 3rd FlipFlop in be 1 and 0 respectively. This asserts RESETn keeping LIT in reset. When the battery voltage rises higher than 1.7V, both the 1.7V and 1.2V Comparators' outputs become high and latches the 2nd Flip Flop from an unknown state (X) to 1 from the output of the 1st Flip Flop. This in turns sets the 3rd Flop Flop making its output 1, deasserting Reset, releasing the system and allowing it to operate. The deassertion of Reset (RESETn becomes 1) causes the RN ports of the 1st and 2nd Flip Flops to force the values of their outputs to 0 which occurs some time after a delay line represented by the inverter. During normal operation, when LIT is operating, the voltage droops a little but never below 1.7V and reset never gets asserted. Once LIT's Carbon-Zinc batteries voltage drop below 1.7V, The 1.7V comparator's output becomes 0, triggering the AND gate to set a 0 on the 3rd Flip Flop's RN, therefore asserting reset. If the battery voltage creeps up past 1.7V due to self-healing, the AND gate asserts a 1. This would have released the system in a more traditional implementation. However, due to the 1 stored at the output of the first Flip Flop, the proposed architecture does not release reset for the system since a 0 is stored at the output of the first Flip Flop. If the battery voltage drops below 1.2V, the 1.2V comparator causes the OR gate to output a 0 as well, resetting the system to its initial state. This means that when the battery goes back up past 1.7V again, the Power-on-Reset / Brown-out-Detector will deassert reset again. Hence, the hysteresis of the system is 1.7V - 1.2V = 500mV. This is sufficient to avoid oscillations due to the self-healing property of the battery and the system is reset only through battery removal which causes the battery voltage to drop to 0 volts.

With the proposed circuit, LIT has a robust Power-on-Reset / Brown-out-Detector. With its unique dual comparators and hysteresis, the Power-on-Reset / Brown-out-Detector has a lock-off feature that overcomes the oscillation due to the self healing properties of Carbon Zinc batteries that circumvents continuously turning on and off that previous works would not have been able to overcome. Using entirely on-chip components, we satisfy the overarching low cost requirement of LIT.

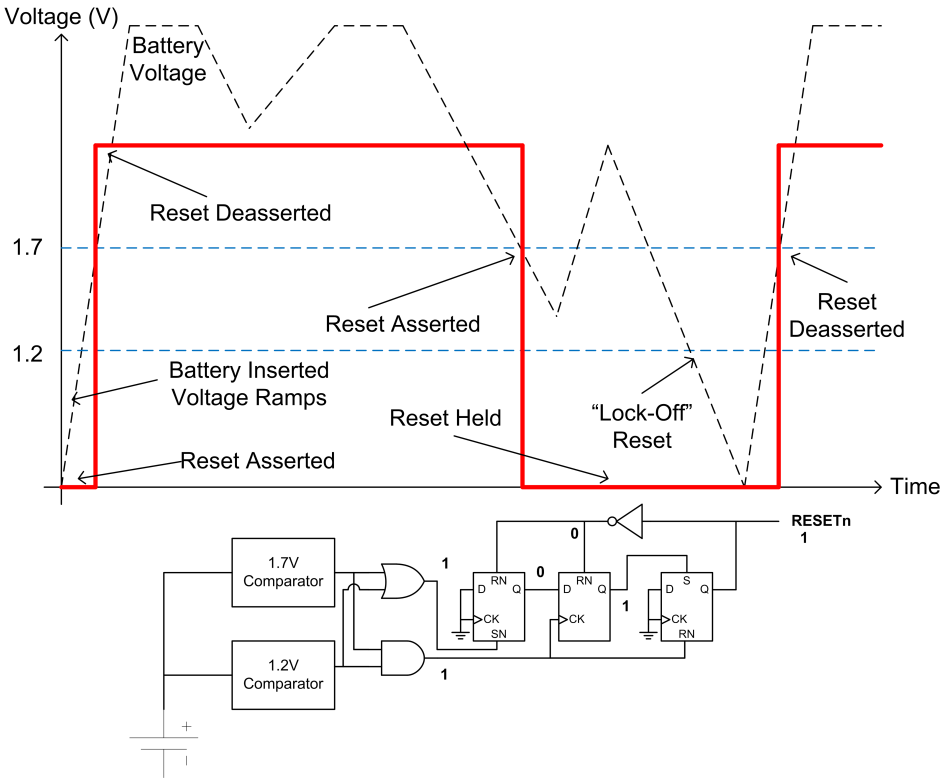

Figure 7.4: LIT's on-chip 1.7μA Power-on-Reset / Brown-out-Detector with "lock-off" operation.

CHAPTER 8

Robustness and Cost Sensitivity through Overloading PCB Traces

8.1 Introduction

In this section we discuss the approaches used in the LIT system for communication to and from the device. This includes both human input through a type of key board and electronic links that allow the device to communicate with other devices. We present our method for both of these interface channels, which uses a novel application specific solution based on the reuse or overloading of PCB traces on the LIT board.

As discussed in Section 2.3, cost and robustness are key objectives and, if possible, we want to satisfy both constraints simultaneously. Particularly in the area of communication channels, the frequent presence of mechanical parts (such as keyboards, cables, etc.) are a liability to robustness. This is due to end users' inexperience in using ICT devices, as discussed in Chapter 2.

For human input interfaces, previous ICT works such as Radios, Cassette Players, Cellular Phones, The Saber [13], Speaking Books [20], One Laptop Per Child [32], and the Talking Book [21], all use some form of push membrane or mechanical switch for its input keypad or keyboard. This simplifies the human input interface from a design complexity's perspective. However, these mechanical parts are a potential failure point for these ICT devices. For instance, the keyboard of the One Laptop Per Child [32] has been known to fail when keys

get stuck [110], or because the keys get torn off by the end users [111] [112]. In addition, these membrane input devices and switches are a substantial cost factor, adding significantly to the total cost of the device.

Data transfer from one device to another also frequently requires some sort of physical port or additional off-chip components. Literacy Bridge's Talking Book [21] uses USB for transfer from one device to another, requiring a separate USB chip, port, and wire, and the power system to power the USB chip itself. However, these ports reduce the robustness of the system since they can be accessed by the end users and could be unintentionally abused or used incorrectly. In order to increase the robustness of these device, these open ports must be removed or kept to a minimum. One Laptop Per Child [32] uses 5GHz Wi-Fi and Bluetooth communication, which does not involve mechanical components. However, it involves additional chips and their associated passives and antennas for those communication protocols as well. These additional chips and their associated passive components run counter to LIT's design goals of reducing the total number of components off-chip in order to reduce its total final cost.

LIT improves upon previous ICT devices by uniquely using the traces on PCBs (Figure 8.1) to implement both human input interface and information transfer from one ICT device to another. By using only traces on the PCB, a number of the mechanical parts found in previous designs are removed. Furthermore, through the integration of all the circuit designs on-chip and not requiring any passives off-chip, the cost of LIT's human input interface and information transfer is significantly reduced since PCB traces are already available on the PCB thus posing no additional cost. Below, we discuss the solution for human input interface and device to device communication in more detail.

8.2 Human Input Interface Using Capacitance to Digital Converters (CDC) with PCB Traces as Buttons

In order for the end users to interact with the device, LIT uses PCB traces as buttons by sensing their capacitance. By using PCB traces as capacitive sensor buttons, we remove

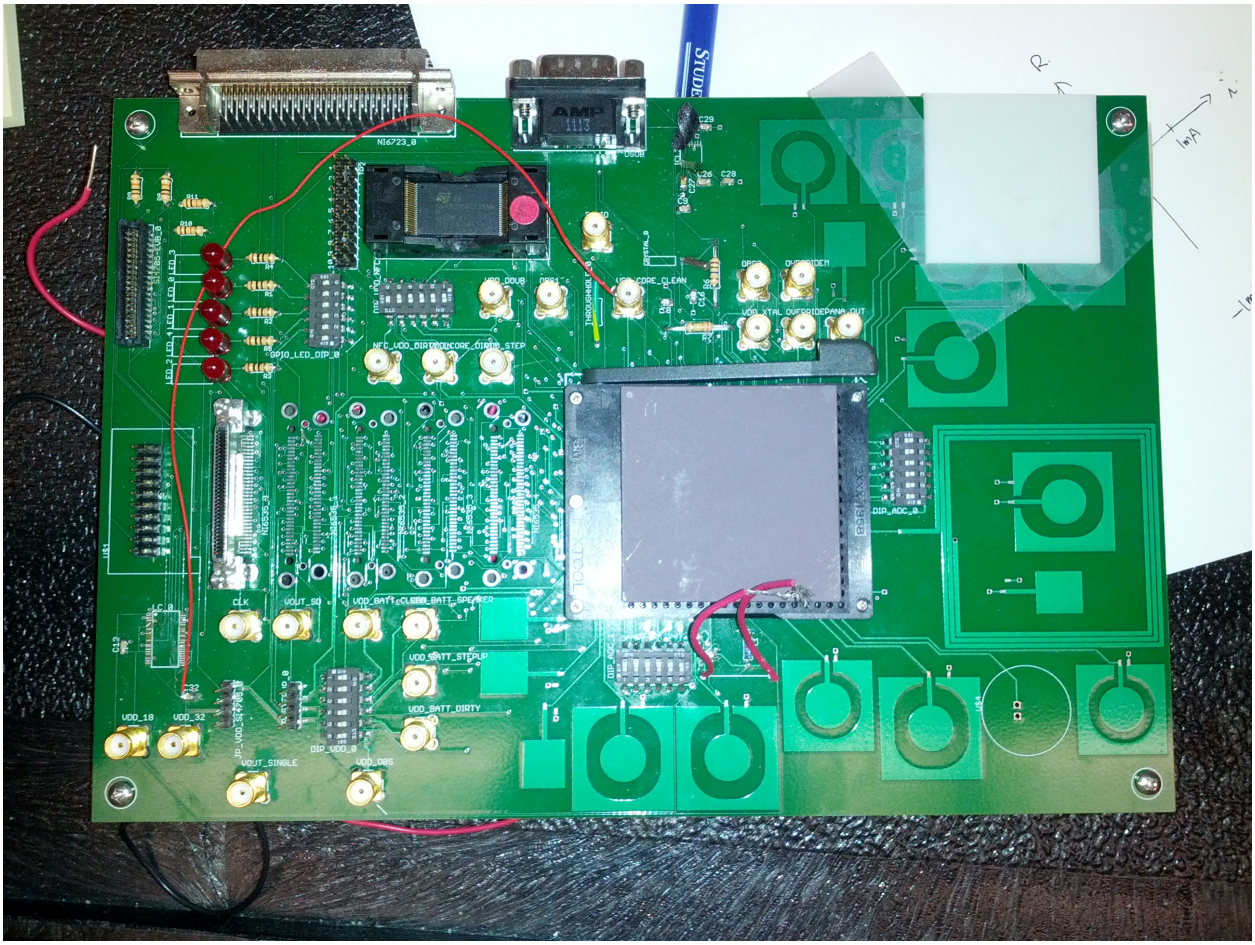

Figure 8.1: LIT's PCB Showing Traces Used as Touch Sensor Human Input Interfaces and as the coil for the inductive link.

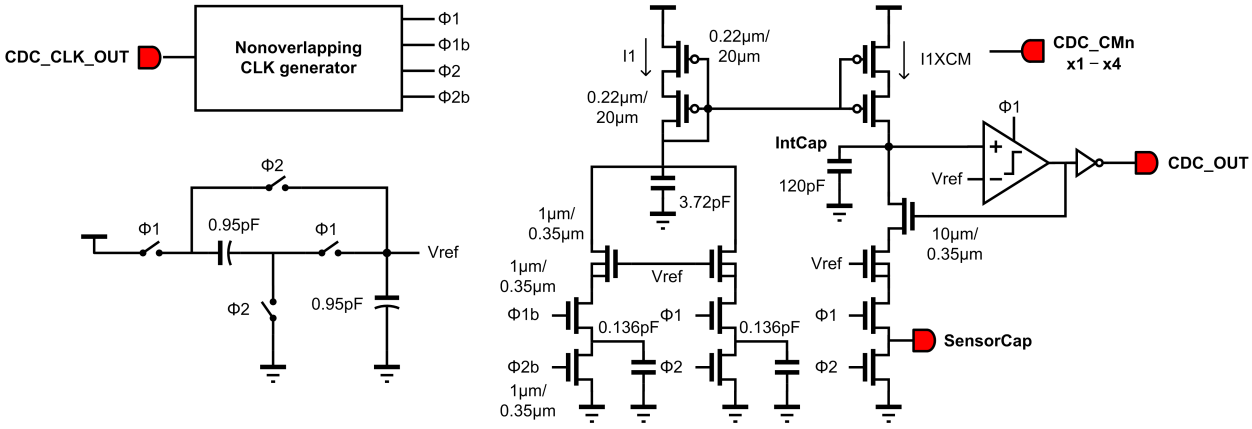

Figure 8.2: LIT's 17-bit Sigma Delta Capacitance to Digital Converter Circuit Level Diagram Requiring only PCB Traces Buttons.

the mechanical element in the human input interface, thereby increasing its robustness. Furthermore, the human input interface can now be entirely enclosed in plastic with no gaps so they cannot get torn off by the end users. Lastly by only using the PCB traces, we reduce the cost of the total device since push membrane buttons would incur a cost overhead for the actual buttons themselves and the interface to the PCB. To enable capacitive buttons, LIT requires a Capacitance to Digital Converter (CDC) in order to sense the buttons. The overhead of the CDC is small since the CDC only requires silicon area and no additional components off-chip.

LIT supports up to 17 touch sensors using the PCB traces as buttons [113] [114] [115] [116] (Figure 8.1) for inputs to a 17-bit Sigma Delta Capacitance to Digital Converter (CDC) [33] [117] [118] [119] [120] [121] (Figure 8.2) (Figure 8.3) designed entirely on-chip for LIT without requiring any further off-chip components. LIT has 10 CDCs with 9 dedicated inputs and another 8 inputs that can be muxed into a single CDC allowing that CDC's functionality to be overloaded in the case that 10 buttons are not enough. This allows LIT to maintain some flexibility in the number of buttons that it can support while maintaining a low area overhead. Although the circuit techniques for the CDC are known, they have not been implemented in conjunction with PCB traces as buttons in currently available ICT devices listed in Chapter 2.

Unlike simple push button membrane human input interfaces that can be simply tied into GPIO inputs typically found in the microcontrollers for such ICT devices, using CDCs

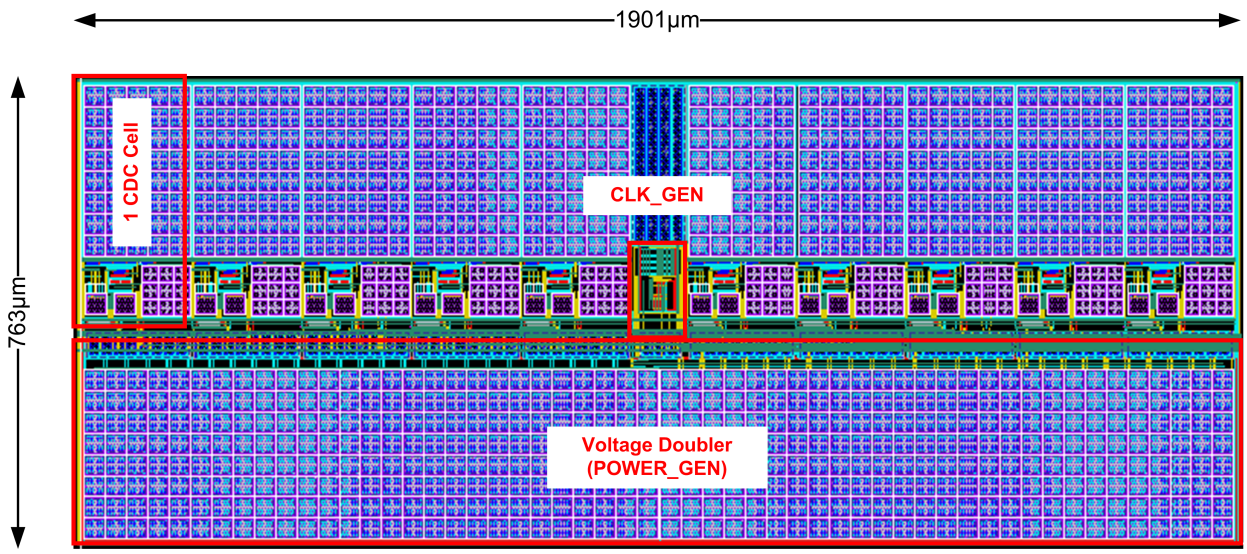

Figure 8.3: LIT's 17-bit Sigma Delta Capacitance to Digital Converter Layout.

as human input interfaces is more complex. There is a possibility of a heavy amount of processing that would need to be done in order to calibrate and monitor the touch sensors for drift in values, debouncing, and determining an actual button event. Given the large code base that would be required to monitor the multiple CDCs and in order to free up the Cortex-M0, LIT employs a hardware based CDC value monitoring circuit that keeps the last 16 values of the CDC outputs and has 2 major ways of determining if there has been valid button activity. The first way is that it merely compares the last output of the CDC and checks it against two thresholds, a positive threshold, or a negative threshold, to see if the value has exceeded either of them. If it has, the CDC's digital Back-End fires an interrupt to the processor to service that particular button. The second way is through a configurable digital low-pass filter that checks to see if the positive or negative delta of the last 2, 4, 8 or 16 values of the CDC outputs exceeds a threshold. This checks to see if the rate of change instead of the absolute value of the output. Similarly, if the rate of change threshold is exceeded, the digital Back-End fires an interrupt to the processor to come and service the button. Furthermore, any one the CDCs can be individually enabled or disabled depending on the context of use, allowing LIT to save energy when not all buttons are required. The CDC can also configure its sampling duration and sampling frequencies in order to operate in a low activity mode which, while is less sensitive and requires a longer time to register a

button press, can be used to conserve energy.

8.3 Information Transfer Using PCB Traces as Inductive Link Coil

An important criteria for the success of the ICT device is the ability for the end users to transfer information between themselves. This ability can support a number of usage scenarios that are best served with peer-to-peer data exchanges. These include sharing individually recorded messages, downloading audio content, usage statistics, and firmware upgrades. Peer-to-peer communications that mirror human inter-contact networks, and allow viral data exchange, are useful because inexpensive and scalable mechanisms for bi-directional communications are rare (recall that in most cases, mobile phones are too expensive). Cassette Players can easily support this through the transfer of tapes from user to user, but we have argued previously that cassettes are expensive, have moving parts and are a possible failure point from a robustness standpoint.

Other ICT solutions, such as Literacy Bridge's Talking Book, [21] supports device to device communication through a USB connector that connects one device to another. However, in order to support this, a separate USB chip, port, wire and power system is required to achieve this form of communication. Furthermore, the open USB port and wire attached to it is a failure point since it can be accidentally misused by the end users. One Laptop Per Child [32] also supports peer-to-peer data exchanges through its Wi-Fi and Bluetooth capabilities, but similarly requires additional chips [122], power management, and off-chip passive components that all add design complexity, and most importantly, cost to the final device since Wi-Fi chips are in the ∼$10 [123] range.

Similar to the human input interface, we can support peer-to-peer communication with no additional passives on the board, no open ports, wires or dongles, with an entirely on-chip transceiver requiring only silicon area thus achieving LIT's goal of a low cost, robust ICT device. Through LIT's transceiver, LIT can support information transfers using triggered exchanges. That is, when two rural audio computers are placed next to each other and

triggered (e.g. by a button push), they exchange the selected content, as well as any system-queued data, like firmware updates. The inductive link provides all these capabilities at a low cost point, improving upon the current ICT solutions that do not have this capability.

The advantage of using an inductive link for communication, rather than a regular wireless communications chip, is the low complexity of the transceiver. It uses a 5cm X 5cm PCB coil trace (Figure 8.1) and can achieve a bitrate of 1.7Mbps at a distance of up to 6.5cm. The transmitter and receiver share one PCB coil to send or receive data which minimizes pad count and reduces the overall system form factor. Based on extensive 3D EM simulations, a 4-turn 5cm x 5cm coil resulted in maximum transmission distance. On the transmit side, data coming from the core is Manchester encoded in hardware which is used to do On-Off-Keying (OOK) modulation. The 40MHz transmit oscillator uses a complimentary cross-coupled structure to maximize negative impedance while avoiding oscillation voltages above supply voltage. The tank capacitance is digitally tuned to tune the oscillator to the exact desired frequency. The tank inductor is used as transmit coil to minimize the number of off-chip inductors and pads and eliminate the need for a power amplifier. On the receive side, the same capacitor array used in the transmitter resonates out the inductor, thus amplifying the coupled voltage. Then, a low-noise RF amplifier with a gain of 10 amplifies the signal further. The amplifier drives a full-bridge schottky rectifier that then extracts the envelope of the incoming RF signal. A low-pass filter attenuates the high-frequency rectifier output jitter. Afterwards, a data slicer converts the signal to full rail which is then Manchester decoded to regenerate the data stream. The reference for the data slicer comes from a 11-bit R-2R DAC that can be adaptively tuned with 1mV accuracy based on incoming signal strength. Figure 8.4 shows the block diagram of the transceiver. Figure 8.5 shows the layout of the transceiver, 700um X 650um, consuming less than 2% of the total area of LIT.

Measurement results of the transceiver using the PCB coil showed that data rates as high as 1.7 Mbps is possible at 6.5cm separation which is higher than similar work in this area [124].

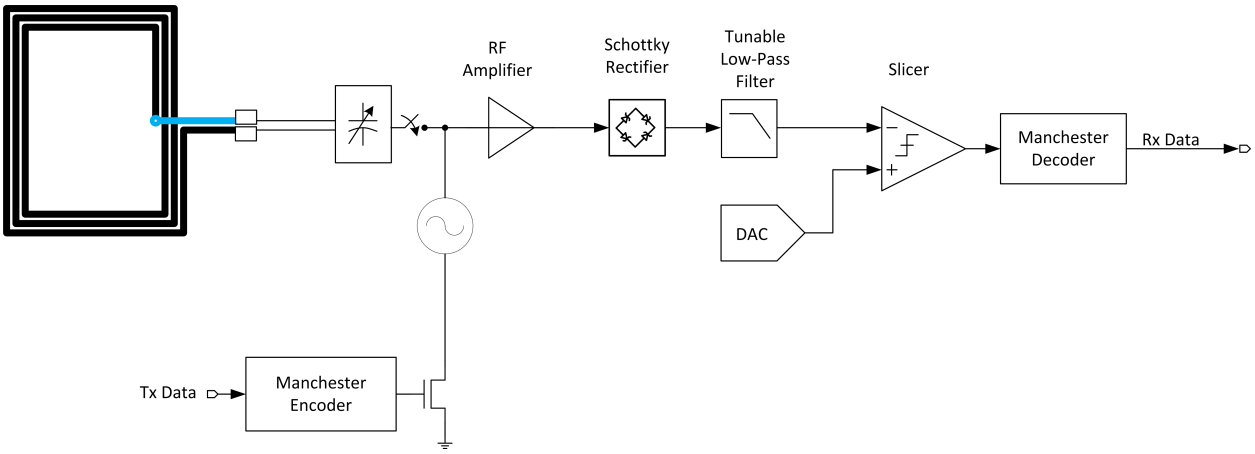

Figure 8.4: LIT's Transceiver for PCB Inductive Link Coil Block Diagram.

8.4 Summary and Novelty of Overloading PCB Traces

LIT shows how by overloading the traces on the PCBs to double as both human input interfaces and a component in data transmission between devices, we manage to achieve a lower cost point and increased robustness when compared to other ICT devices through the elimination of moving parts and open ports. This places a heavy emphasis on the reuse of components already available through the design of a system to lower both cost and improve robustness for systems targeted towards our application area of ICT devices. While inductive communication links and capacitive sensing circuits are not novel in themselves, the realization that they can be applied to the developing world application area and that they hold particular advantages there constitutes the novel contribution of these particular methods.

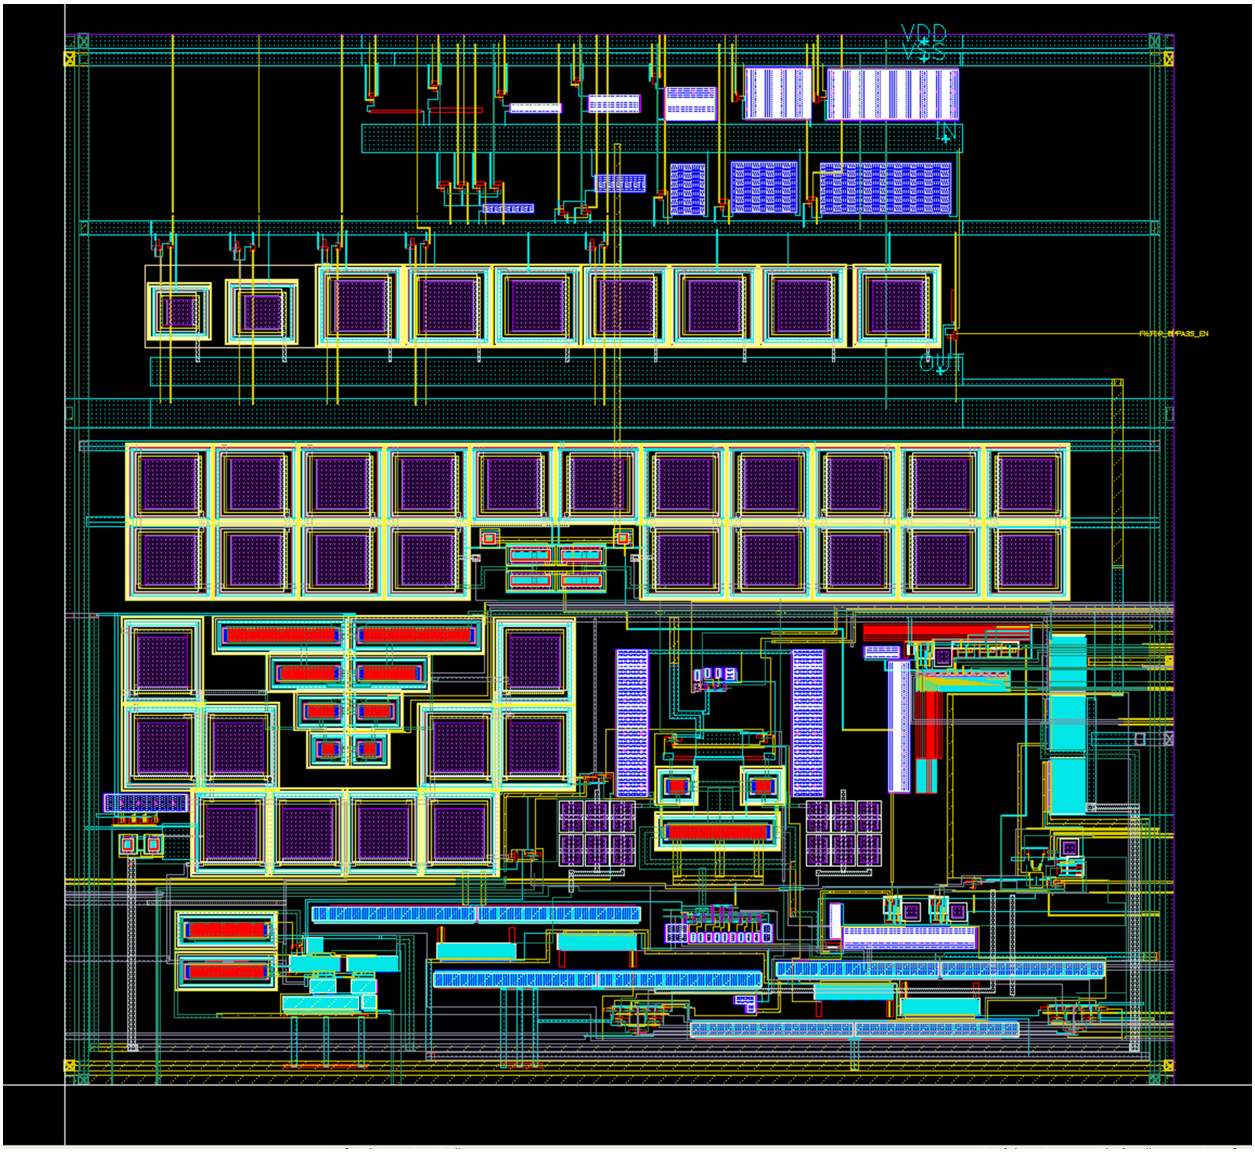

Figure 8.5: LIT's Transceiver Layout for PCB Inductive Link Coil.

CHAPTER 9

Results of Literacy in Technology (LIT), a Low-Cost Custom Silicon Designed for Rural Audio Computers

9.1 LIT's Results

In the previous chapters, we detailed the design decisions that we made in order to meet the design challenges faced that we discussed in Chapter 2. We will now highlight LIT's major results starting with LIT's silicon measured core frequency vs. Voltage and silicon measured core GOPS/WATT (Figure 9.1), Shmoo Plot (Figure 9.2), die micrograph (Figure 9.3), and design parameters and measured results (Table 9.1) shown below.

9.2 LIT's Cost

In Chapter 3, we showed how LIT integrates most components and their associated passives onto a single die and from that, we can significantly reduce part count, pin count, and power draw. This significantly reduces PCB design complexity, which further reduces size, manufacturing, and packaging costs. Table 9.2 shows the bill of materials resulting from this aggressive integration. We expect that the LIT can be produced for less than the target initial cost of $10, a price not met by the other ICT solutions discussed in Chapter 2.

LIT lowers the recurring cost by consuming less energy (3X less) when compared to the current Talking Book (300mW), (5X less) when compared to mobile phones (500mW)

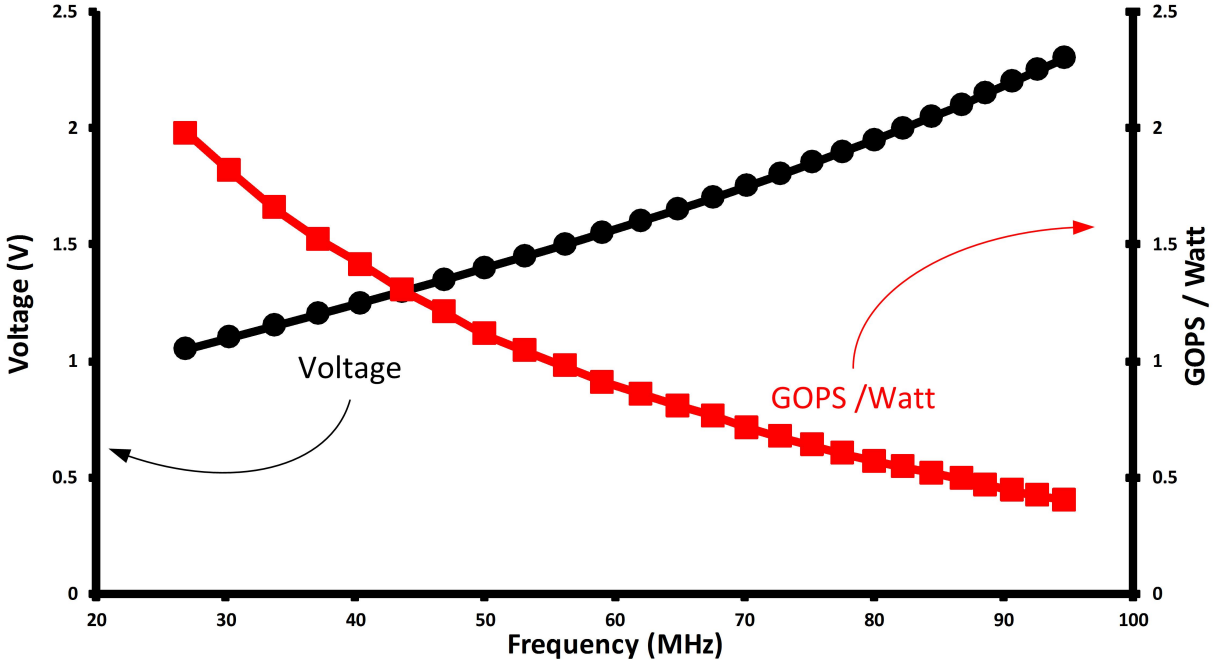

Figure 9.1: LIT's Silicon Measured Core Clock frequency versus Voltage and Silicon Measure Core GOPS/Watt.

available in these areas. Lowering active and sleep energy consumptions are important goals for LIT since this reduces the end users' recurring costs, and so, many of LIT's design choices were made with lowering energy consumption in mind. Clock gating, power gating, and clock speed tuning, depending on workload and module activity, are software configurable and work in unison to lower energy consumption across the board. Furthermore, most of LIT's power hungry components can be power gated individually depending on its use, thereby only using the minimal amount of energy required to do exactly what it's doing at any point in time. LIT's integrated lower energy consumption and low component count allow it to be operated with only two AA batteries for up to two years in Deep Sleep mode, which improves its form factor to a point that makes it more easily transported by end users, currently not possible with the ICT solutions in Chapter 2, and lowers the recurring costs of purchasing new batteries for the device.

In Figure 9.4, we show the cost as a function of production volume. The cost of the custom silicon solution includes the initial mask cost needed to mass-produce the chip, which is approximately \$240k in 0.18μm technology. The relatively low mask set cost in an older

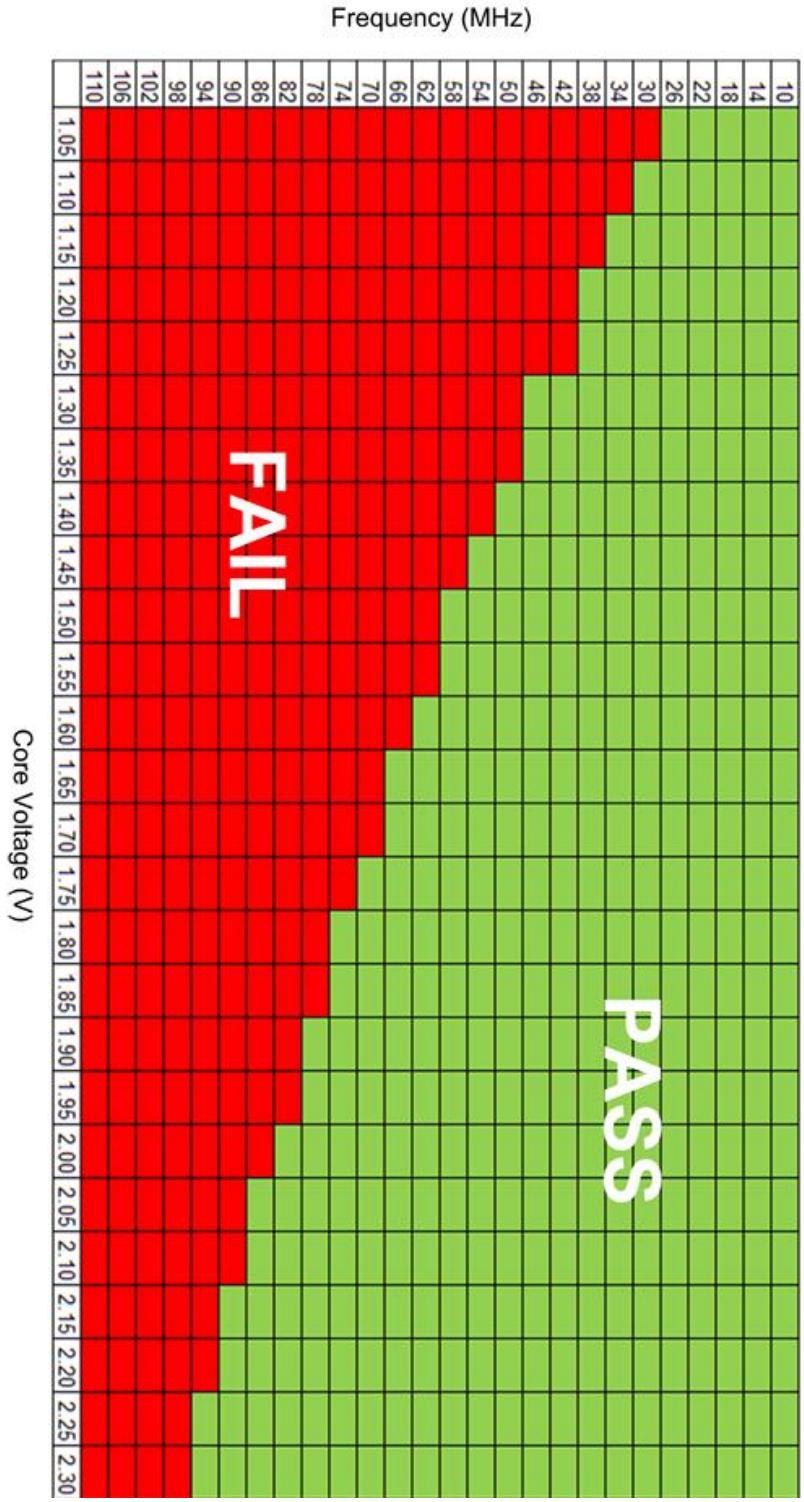

Figure 9.2: LIT's Shmoo Plot.

technologies was a significant factor for choosing the 0.18μm technology for this project. In addition, the design effort was estimated to be eight engineering years at \$160k per year for

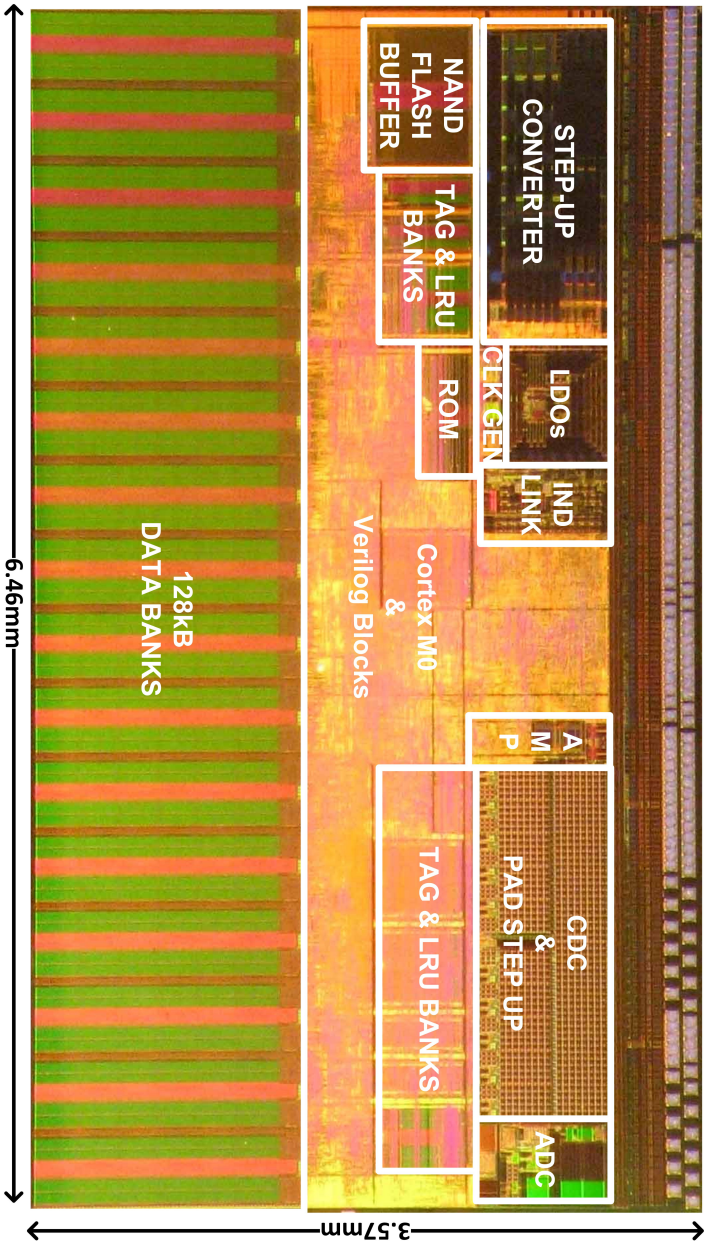

Figure 9.3: LIT's Die Micrograph.

Table 9.1: LIT's design parameters and silicon measured results.

Technology	180nm
Area	23.06mm^2
Performance	64MHz
# Transistors	8Million
# Gates	265K
Cache	128kB True LRU
Cache Area	46% of Total Area
Electronics Cost	<$6
Chip Cost	<$1
ADC **8 DC Input** **4 AC Input** **1 MIC Input**	40KSample/s 3mW 10bit 8.9ENOB 25.6SNDR
Touch Sensor **CDCx10**	3Sample/s 0.72mW 13ENOB
Near Field **Inductive Coil**	BER< 10^{-6} 1.7Mbps at 6.5cm
Active Power	91mW
Standby Power	27mW
Deep Sleep Power	30.6uW
Operation Voltage Range	3.3 — 1.7V
Sleep Lifetime (Carbon Zinc)	2 Years

a total engineering cost of $1.28M. In actuality, the engineering cost was significantly lower since the chip was designed by graduate students at the University of Michigan. Figure 9.4 shows that the cost of both the COTS and custom solution drops as volume increases due to the amortization of non-recurring engineering (NRE) costs and the reduced cost of discrete components at high volumes. However, the COTS solution remains above $20 even for 1M units while the custom solution reaches the required $10 price point at 400k units, motivating the need for a custom solution. The cross-over point between COTS and custom solutions occurs at 70k units, which is relatively small compared to the size of the target market.

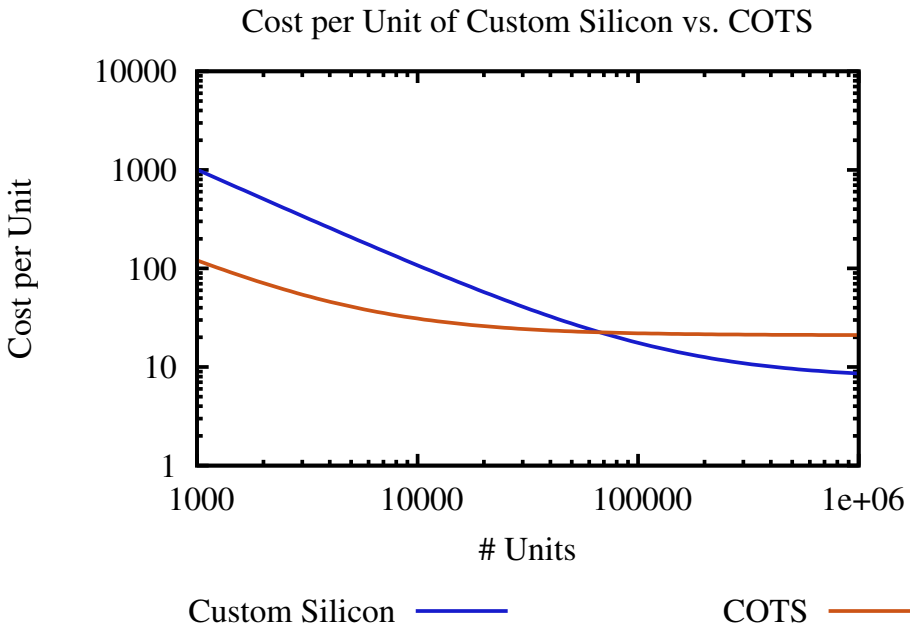

Figure 9.4: Unit cost versus volume for custom silicon and a COTS component design.

9.3 Results Summary

LIT provides a complete and robust power solution for the entire system with a single Carbon Zinc battery input with a wide operating voltage range, slow voltage degradation, and avoiding oscillating voltages near its end life, by integrating the 4 LDOs, Voltage Doubler, Hybrid Switch Capacitor Network, and POR/BOD with lock-off on a single die. It does not require additional off chip components thereby reducing cost and area, fulfilling our design requirements. LIT's low energy consumption through multiple power operational modes and power system designed to exploit Carbon Zinc batteries for efficient operation, novel memory hierarchy that reduces chip count and supports variable power gated banks, coupled with a high level of integration of analog components results in a low board-level component count allowing us to provide a low cost, affordable solution for the dissemination of information to illiterate people groups. The techniques and designs in this thesis allow us to build an ICT solution that meets the challenges of disseminating information among illiterate people groups. We successfully overcome the unique challenges — cost, power, connectivity, robustness, usability, and illiteracy through the design of custom silicon.

Component Name	**Cost**
Chip & Supporting Passives	
Speaker Ind	$0.0245
Speaker Cap	$0.0090
Linear Regulator Cap	$0.0107
Step-Up Converter Cap (Type 1) X4	$0.0027
Step-Up Converter Cap (Type 2) X2	$0.0246
Chip Area	$0.3220
Chip Package	$0.2700
Total Chip Cost	**$0.6635**
Board	
Chip	$0.6635
Speaker	$0.7783
Headphone Jack	$0.0232
Microphone	$0.0935
PCB	$1.0050
NAND Flash	$2.0000
Radio	$1.0395
Total Board Cost	**$5.6031**
System	
Board	$5.6031
Casing	$ 1.0000
Test & Assemble	$1.0000
Total System Cost	**$7.6031**

Table 9.2: LIT's Bill of Materials. Note that the fractional cent costs come from the high volume pricing of components.

CHAPTER 10

Conclusion and Future Directions

10.1 Conclusion

In this thesis, we attempt to improve the quality of life of the poorest people in the world through VLSI technology. We argue that access to information, together with literacy, the ability to read and write, allow people to be independent in their means of absorbing new information and passing of knowledge. Through this, studies have shown that quality of lives can be improved in multiple ways ranging from poverty to child mortality rates. We show that there is a real need to provide timely, actionable information to those who need it the most, the populations that live in rural agrarian settlements, subsistence farmers that are mostly illiterate.

We achieve this by designing custom silicon for a SOC ICT that meets the $6 cost target through a high-level of integration. We also our goals with a novel memory architecture, removal of costly DRAM or NOR Flash, various power operational modes that ensure low energy consumption at any mode of operation through fine granularity of power gating and clock gating, careful design of the Power-on-Reset / Brown-out-Detector catered specifically to this application space with its Carbon Zinc batteries. We also further lowest costs through our Hybrid Switch Capacitor Network with its use of capacitors, power management through the Wakeup Interrupt Controller, increase its robustness, connectivity and usability compared to other ICT designs through the support of multiple communication methods that do not require dongles, physical open ports or wires. By overloading PCB traces used as

both a human input interface device and a component in data transmission between devices, we achieve better robustness and lower cost when compared to current ICT devices.

In conclusion, we presented how Very Large Scale Integrated Chips (VLSI) can improve the quality of life of the poorest people in the world. Through the development of custom silicon designed for an ICT device that can be used to disseminate information to and among the poorest people in the world, identified to be subsistent farmers living in rural agrarian societies in developing regions, we manage to provide timely, actionable information that can improve their quality of life. LIT tackles the unique challenges for this application space — cost, power, connectivity, and robustness.

10.2 General Takeaways

All the solutions derived from LIT have been driven by the need to overcome cost, power, connectivity, usability, robustness, in an ICT device to be used in dissemination information to illiterate agrarian farmers in rural developing regions. Much of the design decisions have been made specifically to cater to the complexities of spreading information in these regions. In the preceding chapters, we have discussed the major technical approaches and insights that were made in the development of the LIT chip. We now look at a number of "takeaways" that were of a more general, non-technical nature, as listed below. These are observations that will carry value, we believe, to other applications in similar developing world projects.

- The unique constraints of developing world applications often result in complex designs with a large number of chips which makes them prohibitively expensive. Chips are primarily manufactured by large semi-conductor companies that target the industrialized world. Hence, we found that they tend to provide a poor fit for the purposes of a developing world application (such as LIT). This can be overcome by combining sufficient different devices together, but this creates a solution with a large component count and high cost. Our first take away is that while it might seem possible to build essentially any type of application given today's selection of COTS chips, the key danger is in creating a large complex system that will end up being too costly. Most ICT devices in the field are highly complex and generally too expensive, meaning they need

subsidies from NGOs and are not economically viable.

- It is key to study the local resources and not make assumptions about which components are locally available. One finding for our device was that modern alkaline batteries are essentially unheard of when one goes to the market places that are in the far corners of developing countries. Similarly, rechargeable batteries, solar panels are also uncommon. Adapting a developed device to the local available resources can significantly increase its usability. In the case of LIT, targeting carbon-zinc batteries had a significant impact on our design decisions.

- In creating a developing world solution, more time is required to determine what to build than is devoted for the construction itself. In this project, a larger fraction of the time was spent on decisions related to what type of chip to build. In the process, it is often necessary to do extensive research into availability of different components and their costs. As an example, in the case of LIT, significant time was invested in understanding the cost and availability of used MicroSD cards, the NAND Flash market, high volume pricing of micro controllers, etc. These types of investigations often drive very important design directions but do not always gain a high-level of visibility in the final description of the project.

- The academic setting poses unique challenges to the successful completion of a large, collaborative project, such as LIT. The complexity of the LIT chip is already such that it is impossible for one student to develop the entire design during his or her thesis work and many students have contributed to the design. In addition, there has been a large team of software developers. Such large collaborative projects are common in industry. However, academia poses a unique challenge in that the training and goals of students are different from that of employees at a company. Utilization of undergraduates for hardware design is challenging since their training level is not sufficient. Senior MS students have sufficient experience to participate in a custom chip design, but they often graduate to quickly to make significant progress. PhD students are available for a longer span of time. However, they must produce a large number of first-author papers. Since a large collaborative project such as LIT produce only a few papers,

the number of PhD students that can be highly invested in it is severely limited. In addition to the technical challenges, the success of the project depends as much on how these human resource issues are managed.

10.3 Future Directions

Moving forward, LIT's second version has been tested and validated at the silicon level. LIT's kernel needs to be finalized with the device drivers for the peripherals validated by the software team which will require some back and forth between the hardware and software teams. After the kernel is finalized, field testing will need to be done. This entails building a stripped down version of the PCB that houses the bare minimum in order to have it portable, operating on batteries but yet still have enough observability to have the ability to debug it in case something goes wrong. There will be considerations taken into account for the PCB design, the size of it, which ports to leave observable, how much hardware overwritability it will have and what sort package we would house LIT in. We will have to carefully design the board to reduce noise on the microphone input and capacitive sensor traces, and where the ground planes will be. We will also reduce the number of layers on the board down to 2 instead of its current 4 layers. We will initially validate it with field testing at University of Michigan by our research team so that we can work out initial bugs with the software and new board configurations. It is possible that field testing could yield further optimization opportunities at the board level. When the prototype is validated to a point where there are no major functional or performance bugs, we will hand it to Literacy Bridge so that they may run their own field tests on top of ours.

We will also have to work with designers to design the plastic housing that will eventually encase the device. There are some design considerations that we have to take into account such as the thickness of the plastic that could potentially affect the sensitivity of the capacitive sensors. We also want to ensure that the device remains robust with the new plastic housing and that the same standards that we have set for LIT apply to it as well; this means no open ports, no dangling wires, and a case that is completely enclosed.

After the plastic housing design is complete, we will start looking for companies that can

produce these at low volume production levels. This means we will have to establish rapports with companies that will produce LIT itself, which we have started doing with Silterra, an alternative TSMC compatible company that produces silicon for smaller companies. We will also have to talk to companies that will do the pick and place for the PCB design, and plastic housing, and eventually a shipping company that will ship the final product to its destinations wherever they may be. We will also have to validate the designs post production from these companies to ensure correct functionality and performance.

We have also established a not-for-profit company called Emergent Microsystems that will be used to talk to these companies and generate relationships with. The IP for LIT will have to be eventually transferred out of Michigan as well and contracts will have to be brought up with University of Michigan's Office of Tech Transfer. Some of the legal issues surrounding tool usage and IP usage from Synopsys and ARM have been solved via 3-way agreements between ARM, University of Michigan and Literacy Bridge ensuring that the design can be taken out into a not-for-profit use scenario.

Finally, further work can be done into where other possible application spaces where LIT can be used. Other example scenarios include populations with defects that impair their vision. There may be changes required that may be able to be accommodated at the board level so that we may reuse LIT. For example, they may not need the inductive coil, or perhaps they can live with a higher initial cost point allowing us to include more features at the board level. We can also potentially accommodate other chips since we support other communications protocols on LIT such as SPI and UART. We also have 27 used GPIO pins so that we can potentially write arbitrary communication protocols with as well which further increases its flexibility.

10.4 Related Works

Z. Foo, D. Devescery, M. Ghaed, I. Lee, A. Madhavan, Y. Park, A. Rao, Z. Renner, N. Roberts, A. Schulman, V. Vinay, M. Wieckowski, D. Yoon, C. Schmidt, T. Schmid, P. Dutta, P. Chen, D. Blaauw., A Case for Custom Dilicon in Enabling Low-cost Information Technology for Developing Regions, *Proceedings of the First ACM Symposium on Computing for Development (ACM DEV 2010)*, Feb 2010.

Z. Foo, D. Devescery, M. Ghaed, I. Lee, A. Madhavan, Y. Park, A. Rao, Z. Renner, N. Roberts, A. Schulman, V. Vinay, M. Wieckowski, D. Yoon, C. Schmidt, T. Schmid, P. Dutta, P. Chen, D. Blaauw. A Low-cost Audio Computer for Information Dissemination among Illiterate People Groups, *Custom Integrated Circuits Conference (CICC)*, Sept. 2012.

Z. Foo, D. Devescery, M. Ghaed, I. Lee, A. Madhavan, Y. Park, A. Rao, Z. Renner, N. Roberts, A. Schulman, V. Vinay, M. Wieckowski, D. Yoon, C. Schmidt, T. Schmid, P. Dutta, P. Chen, D. Blaauw. A Low-cost Audio Computer for Information Dissemination among Illiterate People Groups, *invited paper to Transactions on Circuits and Systems I (TCAS-1)*, June 2013.

BIBLIOGRAPHY

[1] Adult and youth literacy. *United Nations Educational, Scientific and Cultural Organization Institute for Statistics*, September 2012.

[2] Human Development Reports. `http://hdr.undp.org/en/statistics/`.

[3] Neil Patel, Deepti Chittamuru, Anupam Jain, Paresh Dave, and Tapan S. Parikh. Avaaj otalo: a field study of an interactive voice forum for small farmers in rural india. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 733–742, New York, NY, USA, 2010. ACM.

[4] R. Gandhi, R. Veeraraghavan, K. Toyama, and V. Ramprasad. Digital green: Participatory video for agricultural extension. In *Information and Communication Technologies and Development, 2007. ICTD 2007. International Conference on*, pages 1 –10, dec. 2007.

[5] Clifford Schmidt, Trina Jean Gorman, Andrew Azaabanye Bayor, and Michael Shayne Gary. Impact of low-cost, on-demand, information access in a remote ghanaian village. In *Proceedings of the 4th ACM/IEEE International Conference on Information and Communication Technologies and Development*, ICTD '10, pages 43:1–43:10, New York, NY, USA, 2010. ACM.

[6] B.D. Weiss, G. Hart, and R.E. Pust. The relationship between literacy and health. *Journal of Health Care for the Poor and Underserved*, 1(4):351–363, 1991.

[7] P. Sandiford, J. Cassel, M. Montenegro, and G. Sanchez. The impact of women's literacy on child health and its interaction with access to health services. *Population Studies*, 49(1):5–17, 1995.

[8] George T. Bicego and J. Ties Boerma. Maternal education and child survival: A comparative study of survey data from 17 countries. *Social Science & Medicine*, 36(9):1207 – 1227, 1993.

[9] United Nations Education, Scientific and Cultural Organization. Education for all global monitoring report 2010, reaching the marginalized. UNESCO, Paris, France, 2009.

[10] G. Legros, I. Havet, N. Bruce, and S. Bonjour. *The Energy Access Situation in Developing Countries: A Review Focusing on the Least Developed Countries and Sub-Saharan Africa*. UNDP & WHO, New York, 2009.

[11] MobileActive. `http://mobileactive.org/countries`.

[12] `http://www.itu.int/ITU-D/ict/statistics/material/Africa_Village_ICT_2007.pdf`.

[13] The "saber" hand wind digital player. `http://globalrecordings.net/topic/saber/`.

[14] C.K. LaDue, V.V. Sapozhnykov, and K.S. Fienberg. A data modem for gsm voice channel. *Vehicular Technology, IEEE Transactions on*, 57(4):2205–2218, 2008.

[15] J. Sherwani, N. Ali, C.P. Rosé, and R. Rosenfeld. Orality-grounded hcid: Understanding the oral user. *Information Technologies & International Development*, 5(4):pp–37, 2009.

[16] Thomas G. Sticht. Comprehension of repeated time-compressed recordings. *The Journal of Experimental Education*, 37(4):pp. 60–62, 1969.

[17] B. Arons et al. Techniques, perception, and applications of time-compressed speech. pages 169–177, September 1992.

[18] SMS Cost in Africa. `http://manypossibilities.net/2009/04/sms-costs-in-africa-2008/`.

[19] Aditya Dhananjay, Ashlesh Sharma, Michael Paik, Jay Chen, Trishank Karthik Kuppusamy, Jinyang Li, and Lakshminarayanan Subramanian. Hermes: data transmission over unknown voice channels. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, pages 113–124, New York, NY, USA, 2010. ACM.

[20] Speaking & talking book, books that speak. `http://http://www.speakingbooks.com/`.

[21] Literacy bridge, accessible knowledge for all. `http://http://http://www.literacybridge.org/`.

[22] C.K. Prahalad. *The Fortune at the Bottom of the Pyramid*. Wharton School Publishing, New Jersey, 2006.

[23] K. Martin and D. John. *Analog Integrated Circuit Design*. Wiley-IEEE, 1997.

[24] Msp430 microcontroller. `http://www.ti.com/lsds/ti/microcontroller/16-bit_msp430/products.page`.

[25] Silicon labs - energy micro - energy friendly microcontrollers and radios. `http://www.energymicro.com/products/`.

[26] Microcontrollers :: Nxp semiconductors. `http://http://www.nxp.com/products/microcontrollers/`.

[27] Michael Keating and Pierre Bricaud. *Reuse methodology manual: for system-on-a-chip designs*. Springer, 2002.

[28] Resve Saleh, Steve Wilton, Shahriar Mirabbasi, Alan Hu, Mark Greenstreet, Guy Lemieux, Partha Pratim Pande, Cristian Grecu, and Andre Ivanov. System-on-chip: reuse and integration. *Proceedings of the IEEE*, 94(6):1050–1069, 2006.

[29] N Madrid, E Peralías, A Acosta, and Adoración Rueda. Analog/mixed-signal ip modeling for design reuse. In *Proceedings of the conference on Design, automation and test in Europe*, pages 766–767. IEEE Press, 2001.

[30] Han Qi, Zheng Jiang, and Jia Wei. Ip reusable design methodology. In *ASIC, 2001. Proceedings. 4th International Conference on*, pages 756–759. IEEE, 2001.

[31] Zhao Junchao, Chen Weiliang, and Wei Shaojun. Parameterized ip core design. In *ASIC, 2001. Proceedings. 4th International Conference on*, pages 744–747. IEEE, 2001.

[32] One Laptop Per Child. `http://http://one.laptop.org/`.

[33] G. Chen, H. Ghaed, R. Haque, M. Wieckowski, Yejoong Kim, Gyouho Kim, D. Fick, Daeyeon Kim, Mingoo Seok, K. Wise, D. Blaauw, and D. Sylvester. A cubic-millimeter energy-autonomous wireless intraocular pressure monitor. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 310 –312, feb. 2011.

[34] Ye-Sheng Kuo, Sonal Verma, Thomas Schmid, and Prabal Dutta. Hijacking power and bandwidth from the mobile phone's audio interface. In *Proceedings of the First ACM Symposium on Computing for Development*, ACM DEV '10, pages 24:1–24:10, New York, NY, USA, 2010. ACM.

[35] Silicon Laboratories. Broadcast FM radio tuner for consumer electronics. `http://www.silabs.com/pages/DownloadDoc.aspx?FILEURL=Support\%20Documents/TechnicalDocs/Si4704-05-C40.pdf`, December 2009.

[36] G.A. Rincon-Mora and P.E. Allen. A low-voltage, low quiescent current, low drop-out regulator. *Solid-State Circuits, IEEE Journal of*, 33(1):36–44, 1998.

[37] Peter Hazucha, Tanay Karnik, Bradley A Bloechel, Colleen Parsons, David Finan, and Shekhar Borkar. Area-efficient linear regulator with ultra-fast load regulation. *Solid-State Circuits, IEEE Journal of*, 40(4):933–940, 2005.

[38] Ka Nang Leung and Philip KT Mok. A cmos voltage reference based on weighted δv¡ sub¿ gs¡/sub¿ for cmos low-dropout linear regulators. *Solid-State Circuits, IEEE Journal of*, 38(1):146–150, 2003.

[39] Robert J Milliken, Jose Silva-Martínez, and Edgar Sanchez-Sinencio. Full on-chip cmos low-dropout voltage regulator. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 54(9):1879–1890, 2007.

[40] Vishal Gupta, Gabriel A Rincón-Mora, and Prasun Raha. Analysis and design of monolithic, high psr, linear regulators for soc applications. In *SOC Conference, 2004. Proceedings. IEEE International*, pages 311–315. IEEE, 2004.

[41] K.E. Kuijk. A precision reference voltage source. *Solid-State Circuits, IEEE Journal of*, 8(3):222 –226, june 1973.

[42] R.J. Baker. *CMOS Circuit Design, Layout and Simulation*. Wiley-IEEE, 2005.

[43] Mitsuru Hiraki, Kenichi Fukui, and Takayasu Ito. A low-power microcontroller having a 0.5-μa standby current on-chip regulator with dual-reference scheme. *Solid-State Circuits, IEEE Journal of*, 39(4):661–666, 2004.

[44] Linear Regulator (LDO) - Single Channel LDO - TPS78001 - TI.com. `http://www.ti.com/product/tps78001`.

[45] AT91SAM ARM-based Flash MCU. `http://www.atmel.com/images/atmel_6257_32-bit-arm7tdmi-microcontroller_sam7l_datasheet.pdf`.

[46] Digikey AT91SAM ARM-based Flash MCU. `http://www.digikey.com/product-detail/en/AT91SAM7L64-CU/AT91SAM7L64-CU-ND/1886176`.

[47] R.A. Powers. Batteries for low power electronics. *Proceedings of the IEEE*, 83(4):687 –693, apr 1995.

[48] F. Trinidad, MC Montemayor, and E. Fatas. Performance study of zn/zncl2, nh 4cl/polyaniline/carbon battery. *Journal of The Electrochemical Society*, 138(11):3186–3189, 1991.

[49] Kanishka Lahiri, Sujit Dey, Debashis Panigrahi, and Anand Raghunathan. Battery-driven system design: A new frontier in low power design. In *Proceedings of the 2002 Asia and South Pacific Design Automation Conference*, ASP-DAC '02, pages 261–, Washington, DC, USA, 2002. IEEE Computer Society.

[50] Bishop Brock and Karthick Rajamani. Dynamic power management for embedded systems [soc design]. In *SOC Conference, 2003. Proceedings. IEEE International [Systems-on-Chip]*, pages 416–419. IEEE, 2003.

[51] Luca Benini, Alessandro Bogliolo, and Giovanni De Micheli. A survey of design techniques for system-level dynamic power management. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 8(3):299–316, 2000.

[52] Yung-Hsiang Lu and Giovanni De Micheli. Comparing system level power management policies. *Design & Test of Computers, IEEE*, 18(2):10–19, 2001.

[53] Luca Benini and Giovanni de Micheli. System-level power optimization: techniques and tools. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 5(2):115–192, 2000.

[54] Maxim Integrated. White-led boost converter vs. charge pump. `http://pdfserv.maximintegrated.com/en/an/AN3243.pdf`, Jan 2005.

[55] Eric J Carlson, Kai Strunz, and Brian P Otis. A 20 mv input boost converter with efficient digital control for thermoelectric energy harvesting. *Solid-State Circuits, IEEE Journal of*, 45(4):741–750, 2010.

[56] Yogesh K Ramadass and Anantha P Chandrakasan. A battery-less thermoelectric energy harvesting interface circuit with 35 mv startup voltage. *Solid-State Circuits, IEEE Journal of*, 46(1):333–341, 2011.

[57] J-P Im, S-W Wang, S-T Ryu, and G-H Cho. A 40 mv transformer-reuse self-startup boost converter with mppt control for thermoelectric energy harvesting. *Solid-State Circuits, IEEE Journal of*, 47(12):3055–3067, 2012.

[58] P Chen, Koichi Ishida, Katsuyuki Ikeuchi, Xin Zhang, Kentaro Honda, Yasuyuki Okuma, Yoshikatsu Ryu, Makoto Takamiya, and Takayasu Sakurai. A 95mv-startup step-up converter with v¡ inf¿ th¡/inf¿-tuned oscillator by fixed-charge programming and capacitor pass-on scheme. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 216–218. IEEE, 2011.

[59] Po-Wa Lee, Yim-Shu Lee, David KW Cheng, and Xiu-Cheng Liu. Steady-state analysis of an interleaved boost converter with coupled inductors. *Industrial Electronics, IEEE Transactions on*, 47(4):787–795, 2000.

[60] Soumitro Banerjee and Krishnendu Chakrabarty. Nonlinear modeling and bifurcations in the boost converter. *Power Electronics, IEEE Transactions on*, 13(2):252–260, 1998.

[61] Gang Yao, Alian Chen, and Xiangning He. Soft switching circuit for interleaved boost converters. *Power Electronics, IEEE Transactions on*, 22(1):80–86, 2007.

[62] Ju-Won Baek, Myung-Hyo Ryoo, Tae-Jin Kim, Dong-Wook Yoo, and Jong-Soo Kim. High boost converter using voltage multiplier. In *Industrial Electronics Society, 2005. IECON 2005. 31st Annual Conference of IEEE*, pages 6 pp.–, 2005.

[63] Yungtaek Jang and M.M. Jovanovic. New two-inductor boost converter with auxiliary transformer. In *Applied Power Electronics Conference and Exposition, 2002. APEC 2002. Seventeenth Annual IEEE*, volume 2, pages 654–660 vol.2, 2002.

[64] Alberto Reatti and Marian K Kazimierczuk. Small-signal model of pwm converters for discontinuous conduction mode and its application for boost converter. *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on*, 50(1):65–73, 2003.

[65] Chi Yat Leung, Philip KT Mok, and Ka Nang Leung. A 1-v integrated current-mode boost converter in standard 3.3/5-v cmos technologies. *Solid-State Circuits, IEEE Journal of*, 40(11):2265–2274, 2005.

[66] Jonathan HB Deane and DC Hamill. Analysis, simulation and experimental study of chaos in the buck converter. In *IEEE Power Electronics Specialists Conference*, volume 2, pages 491–498, 1990.

[67] Enric Fossas and Gerard Olivar. Study of chaos in the buck converter. *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on*, 43(1):13–25, 1996.

[68] Jinwen Xiao, Angel V Peterchev, Jianhui Zhang, and Seth R Sanders. A 4-μa quiescent-current dual-mode digitally controlled buck converter ic for cellular phone applications. *Solid-State Circuits, IEEE Journal of*, 39(12):2342–2348, 2004.

[69] Kaiwei Yao, Mao Ye, Ming Xu, and Fred C Lee. Tapped-inductor buck converter for high-step-down dc-dc conversion. *Power Electronics, IEEE Transactions on*, 20(4):775–780, 2005.

[70] Richard Redl, Brian P Erisman, and Zoltan Zansky. Optimizing the load transient response of the buck converter. In *Applied Power Electronics Conference and Exposition, 1998. APEC'98. Conference Proceedings 1998., Thirteenth Annual*, volume 1, pages 170–176. IEEE, 1998.

[71] Volkan Kursun, Siva G Narendra, Vivek K De, and Eby G Friedman. Analysis of buck converters for on-chip integration with a dual supply voltage microprocessor. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 11(3):514–522, 2003.

[72] Alejandro R Oliva, Simon S Ang, and Gustavo Eduardo Bortolotto. Digital control of a voltage-mode synchronous buck converter. *Power Electronics, IEEE Transactions on*, 21(1):157–163, 2006.

[73] Vatché Vorpérian. Simplified analysis of pwm converters using model of pwm switch. ii. discontinuous conduction mode. *Aerospace and Electronic Systems, IEEE Transactions on*, 26(3):497–505, 1990.

[74] Gaetano Palumbo and Domenico Pappalardo. Charge pump circuits: An overview on design strategies and topologies. *Circuits and Systems Magazine, IEEE*, 10(1):31–45, 2010.

[75] J.F. Dickson. On-chip high-voltage generation in mnos integrated circuits using an improved voltage multiplier technique. *Solid-State Circuits, IEEE Journal of*, 11(3):374 – 378, jun 1976.

[76] L. Pylarinos. Charge pumps: An overview. In *Proceedings of the IEEE International Symposium on Circuits and Systems*. Citeseer, 2003.

[77] Y.K. Ramadass and A.P. Chandrakasan. Voltage scalable switched capacitor dc-dc converter for ultra-low-power on-chip applications. In *Power Electronics Specialists Conference, 2007. PESC 2007. IEEE*, pages 2353 –2359, june 2007.

[78] E. Racape and J.-M. Daga. A pmos-switch based charge pump, allowing lost cost implementation on a cmos standard process. In *Solid-State Circuits Conference, 2005. ESSCIRC 2005. Proceedings of the 31st European*, pages 77 – 80, sept. 2005.

[79] Pierre Favrat, Philippe Deval, and Michel J Declercq. A high-efficiency cmos voltage doubler. *Solid-State Circuits, IEEE Journal of*, 33(3):410–416, 1998.

[80] Hongchin Lin and Nai-Hsien Chen. New four-phase generation circuits for low-voltage charge pumps. In *Circuits and Systems, 2001. ISCAS 2001. The 2001 IEEE International Symposium on*, volume 1, pages 504–507. IEEE, 2001.

[81] Jieh-Tsorng Wu, Yueh-Huang Chang, and Kuen Long Chang. 1.2 v cmos switched-capacitor circuits. In *Solid-State Circuits Conference, 1996. Digest of Technical Papers. 42nd ISSCC., 1996 IEEE International*, pages 388–389. IEEE, 1996.

[82] Jieh-Tsorng Wu and Kuen-Long Chang. Low supply voltage cmos charge pumps. In *Symp. VLSI Circuits Dig. Tech. Papers*, volume 6, 1997.

[83] Roberto Pelliconi, David Iezzi, Andrea Baroni, Marco Pasotti, and Pier Luigi Rolandi. Power efficient charge pump in deep submicron standard cmos technology. *Solid-State Circuits, IEEE Journal of*, 38(6):1068–1071, 2003.

[84] Toru Tanzawa and Tomoharu Tanaka. A dynamic analysis of the dickson charge pump circuit. *Solid-State Circuits, IEEE Journal of*, 32(8):1231–1240, 1997.

[85] Davide Baderna, Alessandro Cabrini, Guido Torelli, and Marco Pasotti. Efficiency comparison between doubler and dickson charge pumps. In *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on*, pages 1891–1894. IEEE, 2005.

[86] Janusz A Starzyk, Ying-Wei Jan, and Fengjing Qiu. A dc-dc charge pump design based on voltage doublers. *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on*, 48(3):350–359, 2001.

[87] Mohammad R Hoque, Ty McNutt, Jimmy Zhang, Alan Mantooth, and Mohammad Mojarradi. A high voltage dickson charge pump in soi cmos. In *Custom Integrated Circuits Conference, 2003. Proceedings of the IEEE 2003*, pages 493–496. IEEE, 2003.

[88] Msp430 ultra-low power 16-bit mcus - value line - msp430g2001 - ti.com. `http://www.ti.com/product/msp430g2001`.

[89] Lpc1112lvfhn24 :: Nxp semiconductors. `http://www.nxp.com/products/microcontrollers/cortex_m0_m0/LPC1112LVFHN24.html`.

[90] Mixed-signal byte-programmable eprom mcu. `https://www.silabs.com/Support\%20Documents/TechnicalDocs/C8051T60x.pdf`.

[91] Rm42l432 16/32-bit risc flash microcontroller. `http://www.ti.com/lit/ds/symlink/rm42l432.pdf`.

[92] Atsam4s16b. `http://www.atmel.com/devices/SAM4S16B.aspx`.

[93] Robert Strenz. Embedded flash technologies and their applications: status & outlook. In *Electron Devices Meeting (IEDM), 2011 IEEE International*, pages 9–4. IEEE, 2011.

[94] Nand flash 101: An introduction to nand flash and how to design it in to your next product. `http://download.micron.com/pdf/technotes/nand/tn2919.pdf`.

[95] Min She. *Semiconductor flash memory scaling*. PhD thesis, Johns Hopkins University, 2003.

[96] Two flash technologies compared: Nor vs nand. `http://focus.ti.com/pdfs/omap/diskonchipvsnor.pdf`.

[97] Nand vs. nor flash memory technology overview. `http://umcs.maine.edu/~cmeadow/courses/cos335/Toshiba\%20NAND_vs_NOR_Flash_Memory_Technology_Overviewt.pdf`.

[98] Used and recycled electronics, cell phones and more - recellular. `http://www.recellular.com`.

[99] C. Park, J. Seo, D. Seo, S. Kim, and B. Kim. Cost-efficient memory architecture design of nand flash memory embedded systems. In *Computer Design, 2003. Proceedings. 21st International Conference on*, pages 474 – 480, oct. 2003.

[100] Erik G Hallnor and Steven K Reinhardt. A fully associative software-managed cache design. In *ACM SIGARCH Computer Architecture News*, volume 28, pages 107–116. ACM, 2000.

[101] Danut Burdia, G-A Lazar, L-C Lazar, and N-R Buzatu. On-chip power-on reset strategy and i/o power supply detection for vlsi circuits. In *Signals, Circuits and Systems, 2009. ISSCS 2009. International Symposium on*, pages 1–4. IEEE, 2009.

[102] Hung-Wei Chen and Wen-Cheng Yen. A low power and fast wake up circuit. In *Circuits and Systems, 2004. ISCAS'04. Proceedings of the 2004 International Symposium on*, volume 2, pages II–293. IEEE, 2004.

[103] Jin Zhang, Lin Jiang, and Zecang Zeng. Design of a novel power-on-reset circuit based on power supply detector. In *Scalable Computing and Communications; Eighth International Conference on Embedded Computing, 2009. SCALCOM-EMBEDDEDCOM'09. International Conference on*, pages 355–359. IEEE, 2009.

[104] Sanjay Kumar Wadhwa, GK Siddhartha, and Anand Gaurav. Zero steady state current power-on-reset circuit with brown-out detector. In *VLSI Design, 2006. Held jointly with 5th International Conference on Embedded Systems and Design., 19th International Conference on*, pages 6–pp. IEEE, 2006.

[105] Huy-Binh Le, Xuan-Dien Do, Sang-Gug Lee, and Seung-Tak Ryu. A long reset-time power-on reset circuit with brown-out detection capability. *Circuits and Systems II: Express Briefs, IEEE Transactions on*, 58(11):778–782, 2011.

[106] Lai Xinquan, Yu Weixue, Cao Yu, et al. A low quiescent current and reset time adjustable power-on reset circuit. In *ASIC, 2005. ASICON 2005. 6th International Conference On*, volume 2, pages 559–562. IEEE, 2005.

[107] TR Yasuda, Masaaki Yamamoto, and Takafumi Nishi. A power-on reset pulse generator for low voltage applications. In *Circuits and Systems, 2001. ISCAS 2001. The 2001 IEEE International Symposium on*, volume 4, pages 599–601. IEEE, 2001.

[108] Wen-Ta Lee, Ching-Chieh Lin, Yuh-Shyan Hwang, and Jiann-Jong Chen. New design for a reset ic of mobile device. In *Consumer Electronics, 2009. ISCE'09. IEEE 13th International Symposium on*, pages 124–127. IEEE, 2009.

[109] Jianping Guo, Weiwei Shi, Ka Nang Leung, and Chiu Sing Choy. Power-on-reset circuit with power-off auto-discharging path for passive rfid tag ics. In *Circuits and Systems (MWSCAS), 2010 53rd IEEE International Midwest Symposium on*, pages 21–24. IEEE, 2010.

[110] Widespread keyboard failures on olpc's xo-1 - slashdot. `http://it.slashdot.org/story/08/04/20/176255/widespread-keyboard-failures-on-olpcs-xo-1`.

[111] Olpc keyboards literally being ripped apart. `http://www.engadget.com/2008/03/27/olpc-keyboards-literally-being-ripped-apart/`.

[112] forum.laptop.org view topic - replacement keyboards. `http://en.forum.laptop.org/viewtopic.php?f=1324\&t=150834\&p=548824`.

[113] Mark Lee. The art of capacitive touch sensing. *Planet Analog supplement to EE Times*, pages 4–6, 2006.

[114] Susan Pratt. Sensor pcb design guidelines for the ad7142 and ad7143 capacitance-to-digital converters.

[115] Tom Perme. Layout and physical design guidelines for capacitive sensing, 2007.

[116] Ning Jia. Adi capacitance-to-digital converter technology in healthcare applications.

[117] B George and V Jagadeesh Kumar. Switched capacitor triple slope capacitance to digital converter. *IEE Proceedings-Circuits, Devices and Systems*, 153(2):148–152, 2006.

[118] John O'Dowd, Andreas Callanan, Gabriel Banarie, et al. Capacitive sensor interfacing using sigma-delta techniques. In *Sensors, 2005 IEEE*, pages 4–pp. IEEE, 2005.

[119] Markus Bingesser, Teddy Loeliger, Werner Hinn, Johann Hauer, Stefan Mödl, Robert Dorn, and Matthias Völker. Low-noise sigma-delta capacitance-to-digital converter for sub-pf capacitive sensors with integrated dielectric loss measurement. In *Proceedings of the conference on Design, automation and test in Europe*, pages 868–872. ACM, 2008.

[120] Roumen Nojdelov and Stoyan Nihtianov. Capacitive-sensor interface with high accuracy and stability. *Instrumentation and Measurement, IEEE Transactions on*, 58(5):1633–1639, 2009.

[121] Dong-Yong Shin, Hyunjoong Lee, and Suhwan Kim. A delta–sigma interface circuit for capacitive sensors with an automatically calibrated zero point. *Circuits and Systems II: Express Briefs, IEEE Transactions on*, 58(2):90–94, 2011.

[122] Hdg104-wifi 802.11b/g system in package. `http://media.digikey.com/PDF/Data\%20Sheets/H\&D\%20Wireless\%20PDFs/HDG104_Product_Brief.pdf`.

[123] Hdg104-dn-2 h&d wireless ab — hdg104-dn-2tr-nd — digikey. `http://www.digikey.com/product-detail/en/HDG104-DN-2/HDG104-DN-2TR-ND/2567690`.

[124] Dominik Cirmirakis, Dai Jiang, Andreas Demosthenous, Nick Donaldson, and Tim Perkins. A fast passive phase shift keying modulator for inductively coupled implanted medical devices. In *ESSCIRC (ESSCIRC), 2012 Proceedings of the*, pages 301–304. IEEE, 2012.