A Low-cost Audio Computer for Information Dissemination among Illiterate People Groups

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

University of Michigan Ann Arbor, Literacy Bridge

***Abstract*-We present Literacy in Technology (LIT), a low power, low cost audio processor for information dissemination among illiterate people groups in developing regions. The 265K gate, 8 million transistor, 23mm^2, ARM Cortex M0 processor uses a novel memory hierarchy consisting of an on chip 128kB true LRU cache and off-chip NAND Flash. LIT reduces initial acquisition cost through a high-level of integration that results in a low board-level component count. In addition, it also reduces recurring cost through design decisions that lower energy consumption. LIT's multiple power operational modes and power management schemes are specifically designed for efficient operation on Carbon Zinc batteries. These are commonly found in developing regions and allow LIT to be priced at a point that is viable for illiterate people groups in developing regions.**

I. Introduction

Studies have shown that providing information to developing regions can greatly aid their livelihoods by teaching them about health, agriculture and education. The well-known One Laptop per Child (OLPC) initiative aimed to address this by connecting populations in developing regions to the internet [1]. However, this does not address 17% (824M) of the global population who are illiterate [2] (Fig. 1) and among the poorest in the world. This indicates that there is a serious need for work to be done in order to combat illiteracy. Several other programs have been developed such as "Speaking Books" from Books of Hope which is affordable but is limited to 5-10 minutes of data which cannot be easily updated [3]. Global Recording's Saber's data can be updated but only through a computer and does not have a built-in microphone. Furthermore, it is priced at a point ($45-$65) that is not affordable by the end users [4].

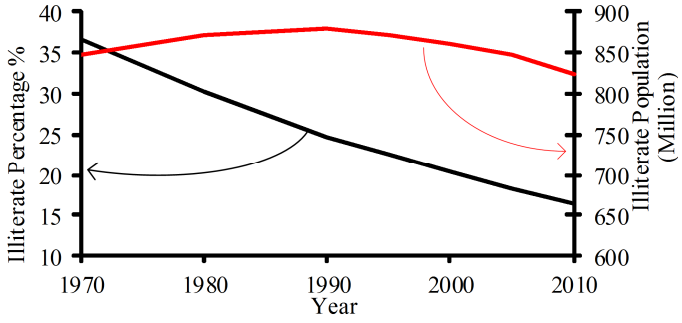

Fig. 1. Worldwide Literacy Rate: Percentage of illiterate people is decreasing, but total illiterate population is remaining constant

The Talking Book (Fig. 2), developed by Literacy Bridge, is a computer with an audio and tactile interface [5]. Using only buttons, 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 [6]. 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.

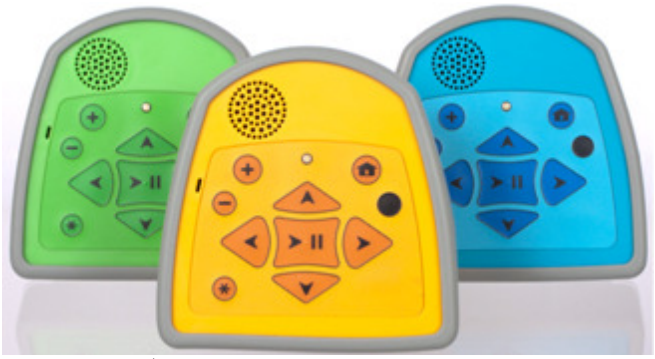

Fig 2. $25 Talking Book developed by Literacy Bridge

The results of the pilot study, collected in January 2010, showed an increase in crop production by an average by 48% when compared to non-users decrease of 5%. 75% of the 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).

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

Given the success of Talking Book's pilot study, its chief obstacle to widely distributing it is to make it economically viable. Its current cost of $25-35 is not affordable by end

users. In contrast, radios typically cost $10 while simple cellphones cost $20. This indicates that the end users can afford a device in the $10 – 20 range. The chief contributor to the current cost of the Talking Books is its 152 off-the-shelf components (Fig. 3). The immediate problem is that the sheer number of off-the-shelf components carries 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 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, compared to a single layer board of 70mm X 100mm in size for the proposed device. While the final device size of the proposed audio computer will eventually depend on the design of the casing, the preliminary board design is 2X smaller, and has fewer layers, thereby reducing PCB cost by more than 80%. Its size can be further reduced as we optimize the component placement and board routing, which will be addressed concurrently with the final casing's design. Finally, the current design's high number of components also affects the size of the plastic housing and shipping costs incurred 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.

Fig 3. Talking Book's PCB and 152 off-the-shelf components

To bridge the cost gap, we developed an SOC which forms the heart of an audio computer for information dissemination among illiterate people groups and has a currently untapped market size of $10B [7]. Using only buttons, the audio computer allows users to navigate through menus and access information which is played back 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 a tactile and audio feedback user-interface, the user can configure and operate the device, in a similar way that made the Talking Book's pilot study a success.

As previously mentioned, the chief obstacle to disseminating these devices in an economically viable manner is the low purchasing power of the end users that typically subsist on $1-2/day. Hence we aim for a total device cost between $10–15, which means all the electronics, including PCB, speakers, microphone, radio, buttons, etc. must remain below $6.

II. Proposed Solution

A. *Overview*

The Literacy Information Technology (LIT) chip (Fig. 4), tackles cost in a number of ways. A novel memory architecture with 128kB 4-way True LRU cache that is directly backed by NAND Flash and removes the need for costly DRAM or NOR Flash commonly found in micro-controllers and processors is used. In addition, LIT integrates all analog components on-chip, including microphone amplifier and biasing circuit, ADC, class-D amplifier, 4 LDOs, 2 voltage references, and 2 step-up converters. The total number of chips in the system is reduced to only 3 (The LIT chip, the NAND Flash chip, and a radio chip), saving cost and PCB size, which in turn reduces the plastic and handling costs (the previously mentioned indirect costs) of the final device. For remaining components, we have made design choices that reduce cost. For instance, the 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. Also, 10 CDCs 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. The resulting 9X reduction in components directly reduces both direct and indirect costs.

While most of LIT's cost reduction was achieved through a reduction of components and substitutions at the board level, we also made a number of design decisions that lowered the cost of the passive devices. For example, instead of using a Step-Up Boost Converter with an expensive inductor for voltage boosting, we developed a new variable capacitive charge pump which cut the cost by 3X by using cheap capacitors.

LIT lowers the recurring cost by consuming less energy (~3X less) when compared to the current Talking Book (~300mW). Lowering active and sleep energy consumptions are important goals for LIT since this reduces the end users' recurring costs. Lowering both active and sleep energy was designed for a targeted application, 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

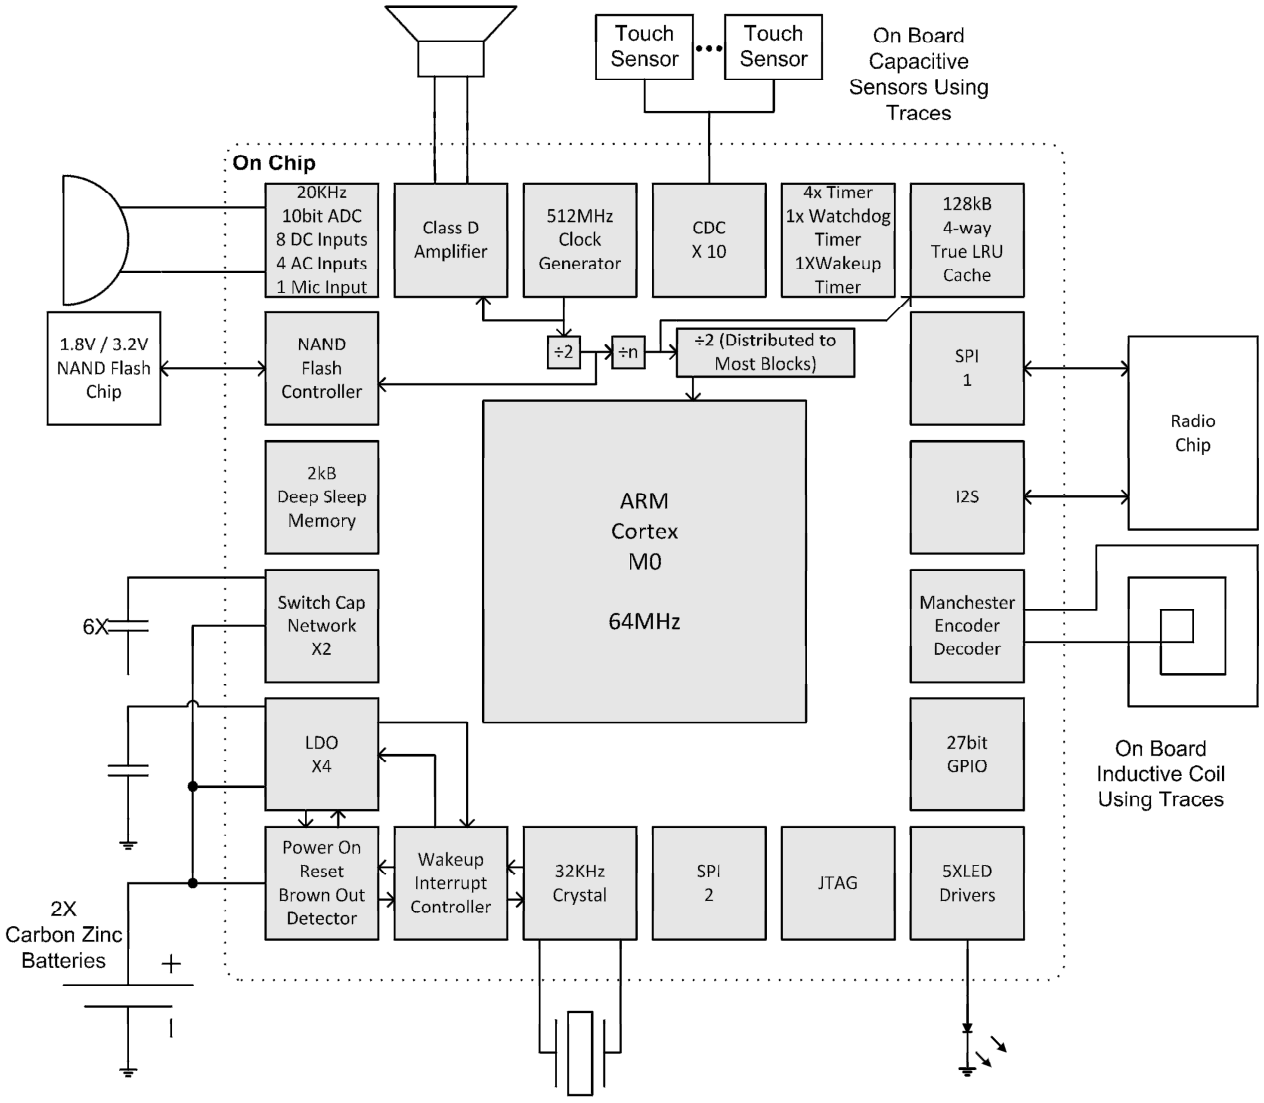

Fig. 4. System Level Diagram Showing On-Chip Components and On-Board Components

configurable and work in unison to lower energy consumption across the board. LIT's integrated lower energy consumption and low component count allows it to be operated with 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 the end users and lowers the recurring costs of purchasing new batteries for the device.

Finally, 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 (Fig. 3) 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. While robustness may not be as large a concern as cost, Talking Book's pilot study showed that end users disassembled the device and left it in the mud/rain indicating that end users may not be as careful with the device as typical users that are more accustomed to electronic technology.

These cost reduction methods produces low component count and minimal board routing (Fig 4), resulting in a small PCB area footprint making it affordable by the end users. The final cost of the system ($7.5), shown in Table 1, meets the target goal of ~$10. The cost breakdown is shown at the Chip, Board, and System levels. We can see that the low number of components in the BOM is a result of our cost reduction methods. Below, we discuss each of the key components of the LIT chip in more detail.

B. Memory Hierarchy

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 (Fig. 4). LIT's memory hierarchy is designed to accommodate the need for large code space (up to 16MB) while eliminating the need for expensive and power hungry NOR Flash or DRAM common in micro-controllers and processors for code

storage. An indirect cost impact of this decision is that LIT has a reduction in the number of pins which reduces testing time as well as packaging costs. LIT implemented a large 128kB 4-way true LRU on-chip cache (48% of die area) directly backed by an off-chip NAND Flash (Fig. 5). The size of the cache is required to reduce the latency times between the NAND Flash and LIT which is not fast enough to support fetching instructions directly. Despite the cache consuming more than 11mm^2 of area, the additional area only increases each die's cost by $0.25, which is still a significant cost advantage when compared to NOR Flash /DRAM, which usually costs $1.

Component Name	Cost
Chip	
Speaker Inductor	$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	$1.0000 (est.)
NAND Flash	$2.0000
Radio	$1.0395
Total Board Cost	**$5.5980**
System	
Board	$5.5980
Casing	$1.0000(est.)
Test & Assemble	$1.0000 (est.)
Total System Cost	**$7.5980**

Table 1. LIT's bill of materials. Note that fractional cent costs come from the high volume pricing of components

LIT uses NAND Flash instead of MicroSD that has overhead in terms of power, access latencies, and most importantly, cost. The MicroSD card accounted for 30% of the cost of the Talking Book's bill of materials. A caveat of using NAND Flash is that we now need to implement our own Flash Translation Layer (FTL) in software. However, by handling cache misses with software, we maintain compatibility with the software FTL, flexibility in prefetching data, and performance through non-blocking cache misses.

The cache has a configurable pinned section that prevents the lower portion of the address space from being evicted which is used for the OS and FTL. This is achieved by having each cache access checked against the address pin line 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 NAND Flash into the cache. The pinned section is required 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.

A consequence of the pinned region is that we must use true LRU 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 is the second least recently used address is in order to evict it, and pseudo LRU does not provide this. Pseudo LRU only provides the least recently used address and nothing else, hence 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 provides as well.

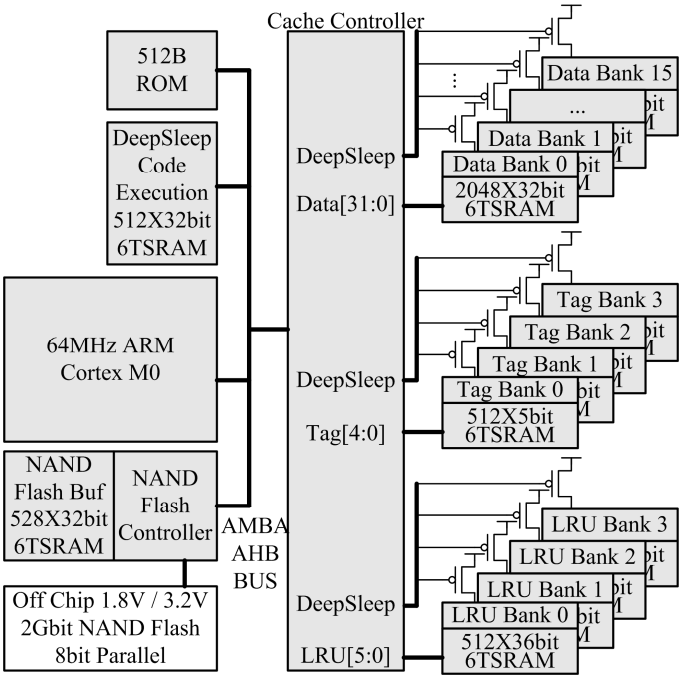

Fig. 5. Memory System Diagram

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. Because of 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.

Upon LIT powering up, a 512B boot ROM 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 sped up through configurations stored in the NAND Flash. This allows LIT to be more flexible and robust. After boot,

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. 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 and when a button is pressed. The event validity checker serves two purposes. First, it is used to validate the requests from the capacitive sensors that we use as the tactile interface since false positives could wake the core unnecessarily, waste energy, and produce unwanted behavior. 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 false positives and 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 Fig. 5). In this case, LIT power gates the entire cache and saves the most energy. Alternatively, we can store the event validity checker in subsections of the cache that is purposely left awake if we require more complicated code that is unable to fit in the 2kB of SRAM. In this option, we are capable of maintaining a variable amount of memory awake 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.

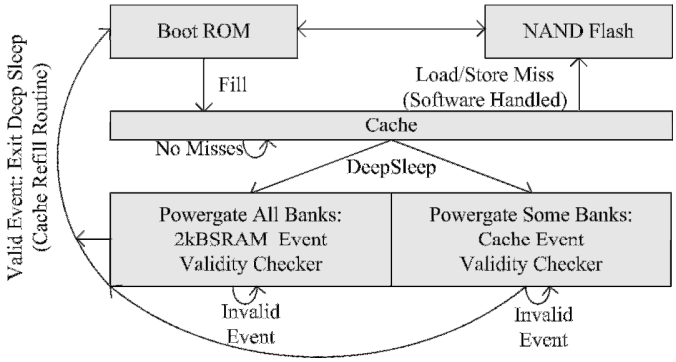

Fig. 6. Memory operation from loading from boot, miss handling during Active mode, and Deep Sleep mode cache power gating options or using 2kB SRAM where event validity checker code is stored.

The cache draws an additional, silicon measured, 1uA in Deep Sleep if 8kB of memory is left awake and 7.2uA if all 128kB is left awake (Table 2). The disproportionally higher current draw of having the first 8kB left awake is due to the need for all LRU & Tag banks to be left awake during Deep Sleep mode to ensure that we know what data is valid. Each additional 8kB of memory left awake during Deep Sleep mode results in an additional 0.4uA in current draw. This indicates that the LRU & Tag banks draw a total of 0.6uA. Having the 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 sleep mode.

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

Table 2. Silicon Measured memory current during Deep Sleep mode

C. *Power Operational Modes*

LIT was designed to have 3 operational modes: Active mode, Standby mode, and Deep Sleep mode in order to reduce energy consumption. It was designed to maximize lifetime when powered by Carbon Zinc batteries, prevalent in developing regions due to their superiority over Alkaline batteries in terms of longer shelf life and lower acquisition cost. Carbon Zinc batteries tend to degrade slower over a larger voltage range unlike Alkaline batteries which have a steep cutoff voltage at 2.2V (Fig. 7). In order to extract maximum energy, it is critical to extract energy from the battery over the entire voltage degradation range. Therefore, we use low drop out regulators and 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 thereby extending the lifetime of the device and lowering recurring cost.

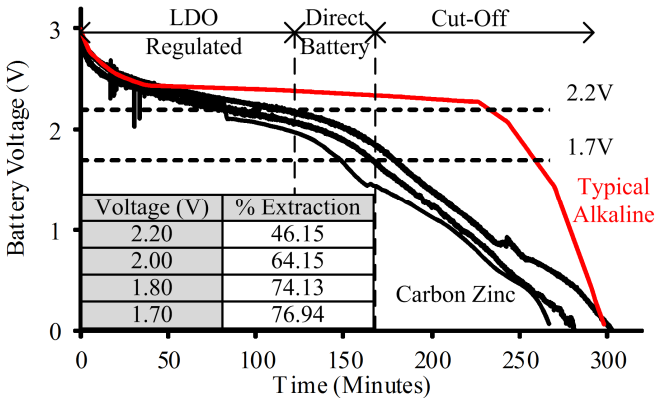

Fig. 7. Measurements of Ghana's Carbon Zinc batteries and typical Alkaline battery. By extending operational voltage from 2.2V to 1.7V, 30% more energy is extracted.

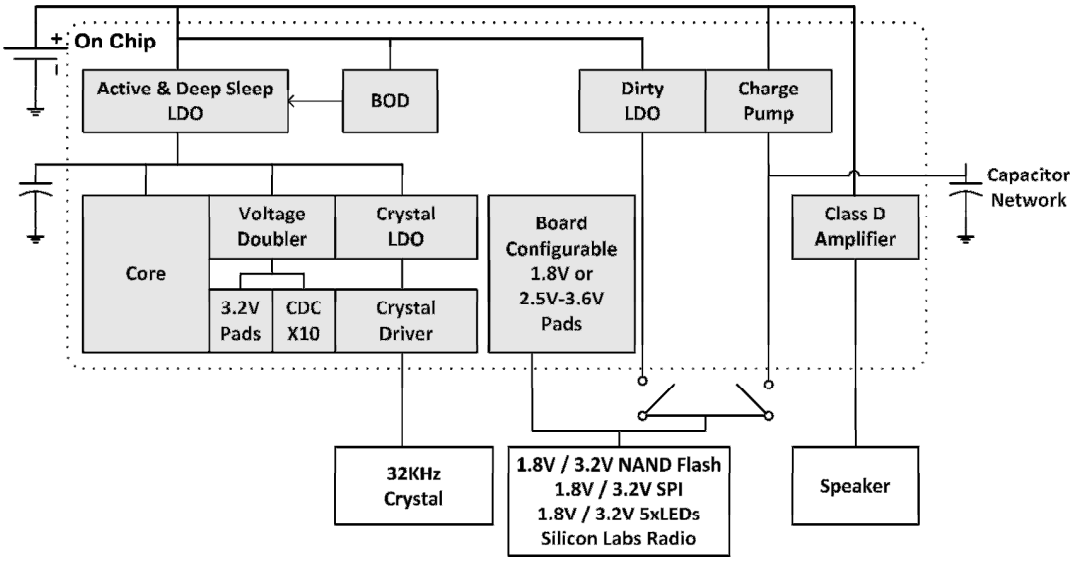

Fig. 9. System Power Diagram: The NAND Flash Controller IO can be powered by either the Variable Charge Pump or the Dirty LDO depending whether a 3.2V or 1.8V NAND Flash is used. We accomplish this by connecting the Variable Charge Pump or Dirty LDO to the PAD Dirty VDD

D. Power Management

In total, LIT has 4 LDOs (Fig 9) with their associated VREFs and IREFs (Fig. 10). The LDOs' various modes are controlled by the Wakeup Interrupt Controller (Section G). Since LIT is sensitive to energy consumption, LIT has its Active mode and Deep Sleep mode each have their own dedicated LDO which results in more efficient energy consumption. During Active mode, the Active LDO is enabled, allowing LIT to draw more current and consume more energy for its workload (mW range). However, when LIT switches to Deep Sleep mode, LIT's energy consumption decreases by ~1000X (uW range), so the Active LDO is power gated and the Deep Sleep LDO takes over regulating the core voltage. 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. The Deep Sleep LDO is continuously enabled during both Active and Deep Sleep mode. During Active mode, the Active LDO is enabled, overpowering the Deep Sleep LDO. This allows for an easier handoff during switching between Active and Deep Sleep modes. Having dedicated LDOs for LIT's different modes allows us to achieve high energy efficiency and to maintain a low Deep Sleep mode energy consumption that provides us up to 2 years of lifetime.

LIT's memory hierarchy requires us to have an on-board NAND Flash chip, and while NAND Flash is cheaper than other solutions (Section B), LIT is flexible 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 a Hybrid Switch Cap Network (Section F), while powering the 1.8V NAND Flash uses LIT's 3rd LDO, the Dirty LDO. The Dirty LDO is similarly controlled by the Wakeup Interrupt Controller (Section G), where it can be enabled or disabled depending on whether we have external components that require 1.8V such as the NAND Flash chip. Having this flexibility lets 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 uses a 32KHz crystal on the board 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 our 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 on board crystal, LIT's 4th LDO, the Crystal LDO (Fig. 10), constantly regulates its 0.95V supply. The Crystal LDO is also constantly enabled since the Wakeup Interrupt Controller's clock operates in the 32KHz clock domain and must always be available. Three of LIT's LDOs, the Active, Deep Sleep and Dirty LDOs use the two phase regulation described in Section C 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.

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 Capacitance-to-Digital-Converters, that both operate at 3.2V. In order to supply the 3.2V, LIT has an on-chip Voltage Doubler (Fig. 10) 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.

Fig. 8. IO Pad Break for Board Voltage Configurable Pads

Unlike most general purpose processors, LIT's LDOs are on-chip thereby reducing the total number of external components which lowers 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 through software for communication with the 1.8V or 3.2V NAND Flash, for example. We overcome this by allowing the IO voltage to be configured at the board level (Fig. 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 (Fig. 8).

E. Power-On-Reset Brown-Out-Detector

In traditional systems, when the battery voltage drops below the system's operating voltage, standard Power-On-Resets/Brown-Out-Detectors (POR/BOD) will assert reset to avoid incorrect operation and memory corruption. However, Carbon Zinc batteries' present a unique challenge in this respect in that the voltage of the battery will start creeping up after reset is asserted due to significantly less current being drawn and the self-healing properties of this battery chemistry. This increased voltage will then reach a voltage where the POR/BOD de-asserts reset and enables the chip thus consuming more current thereby driving the voltage low again. This phenomenon (Fig. 11) leads to an oscillation that will continuously turn the device on and off near the end life of the battery due to the strong self-healing properties of Carbon Zinc batteries, which is clearly unwanted behavior and can result in NAND Flash corruption resulting in loss of data.

In order to circumvent this oscillation, we designed a 1.7uA POR/BOD with a lock-off feature (Fig. 12). In the event of the battery voltage dropping below the BOD low-threshold (1.7V), the device locks off, until the voltage drops significantly further below 1.2V, resetting the lock (i.e, when changing batteries). At this point, if the battery voltage exceeds 1.7V, the device is allowed to be turned on again. This results in a more robust device that is less susceptible to data corruption.

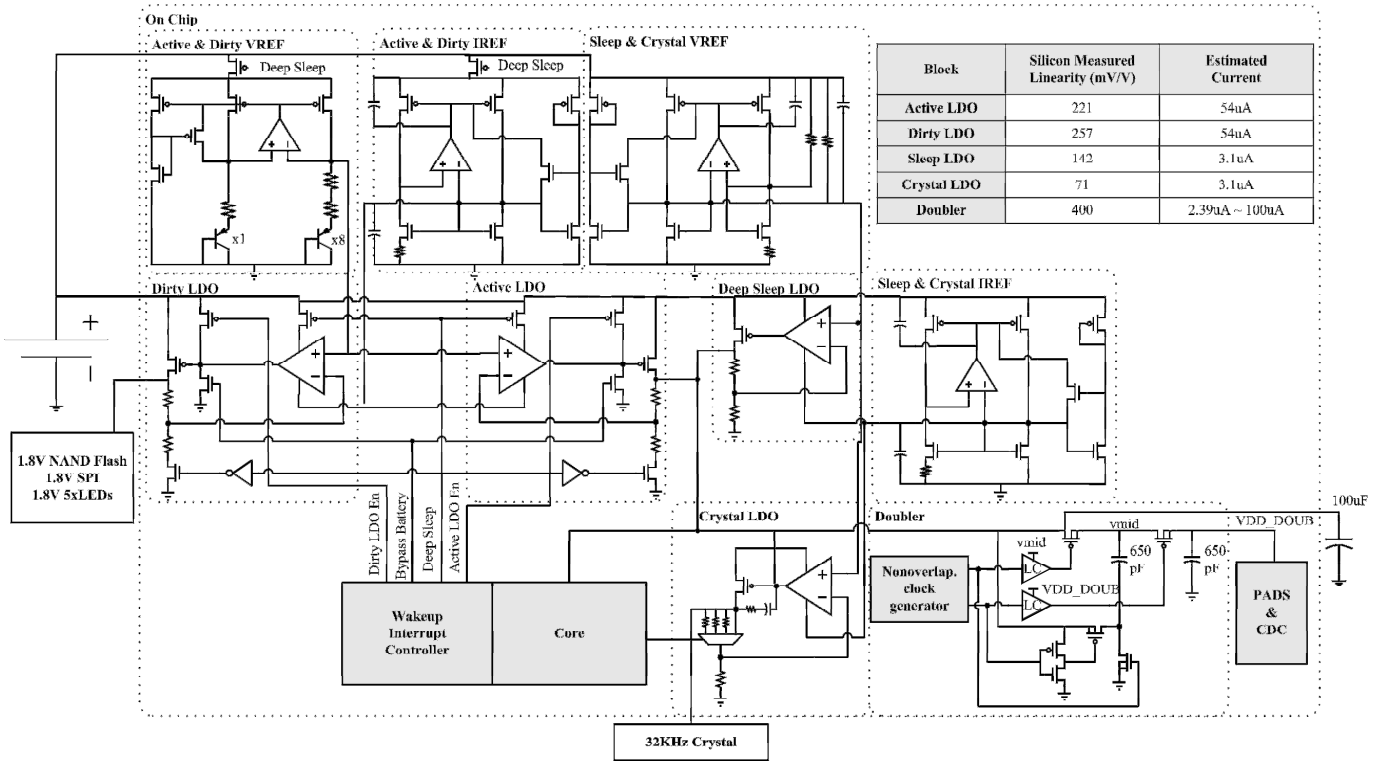

Block	Silicon Measured Linearity (mV/V)	Estimated Current
Active LDO	221	54uA
Dirty LDO	257	54uA
Sleep LDO	142	3.1uA
Crystal LDO	71	3.1uA
Doubler	400	2.39uA ~ 100uA

Fig. 10. LIT's power diagram showing LDOs, and on chip voltage Doubler with their associated VREFs and IREFs and silicon measured output.

Upon battery insertion, the supply voltage is low enough such that both the High and Low Comparators' outputs are low, making the outputs of the 1st and 3rd FlipFlop in the "Lock-Off" be 1 and 0 respectively. This asserts RESETn keeping LIT in reset. When the battery voltage rises higher than 1.7V, both the High and Low Comparators' outputs become high and latches the 2nd Flip Flop from an unknown state to 1 from the output of the 1st Flip Flop. Simultaneously, the 3rd Flip Flop's Reset is deasserted by the High and Low Comparators' output becoming high, and the 2nd Flip Flop's output sets the 3rd Flip Flop to a 1 thus asserting "RESETn" to a 1. This releases reset on LIT allowing it to start fetching its first instruction. The deassertion of "RESETn" to a 1 asserts the 1st and 2nd Flip Flop's Rn thus setting their outputs to 0.

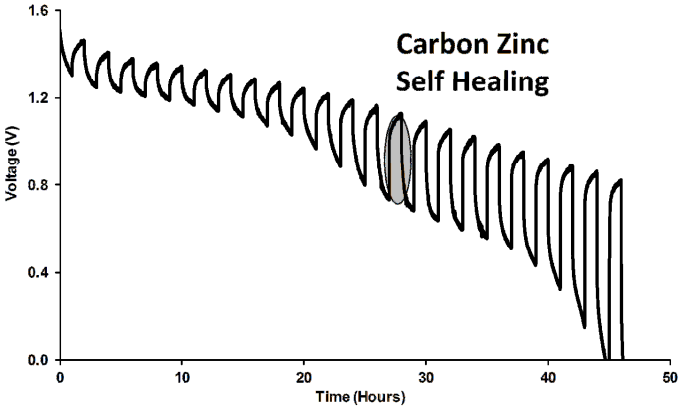

Fig. 11. Carbon Zinc Batteries self healing property

When the battery voltage drops below 1.7V, The High Comparator's output becomes low, asserting the 3rd Flip Flop's Rn, thus making RESETn 0, resetting LIT. In the event that the battery's voltage creeps past 1.7V as described previously, The High Comparator's output becomes high again, but RESETn can never be deasserted because the outputs of the 1st and 2nd Flip Flops in the "Lock-Off" are 0, uneable to assert S on the 3rd Flip Flop. This overcomes the oscillation due to the self healing properties of Carbon Zinc batteries.

In the event that the battery voltage drops below 1.2V, (ie: battery removal) and goes past 1.7V again, (ie: battery insertion), the process repeats.

F. Hybrid Switch Cap Network

Finally, LIT has a second voltage booster to power other chips on the board, such as a 3.2V NAND Flash chip and radio chip. To ensure an output voltage within the 3.2V spec over a wide range of battery voltages (3.2V – 1.7V), most systems opt to use a DC-DC boost converter. While a standard inductive boost-converter would have sufficed to boost the voltage from the battery to the levels that we require, LIT avoids the expense of inductors by creating a new Hybrid Switched Capacitor based boost converter topology. The cost of a single inductor alone is ~$0.38, compared to all the capacitors' cost ~$0.03, that our hybrid Switched Capacitor Network (SCN) boost converter uses, and is therefore the more economical choice.

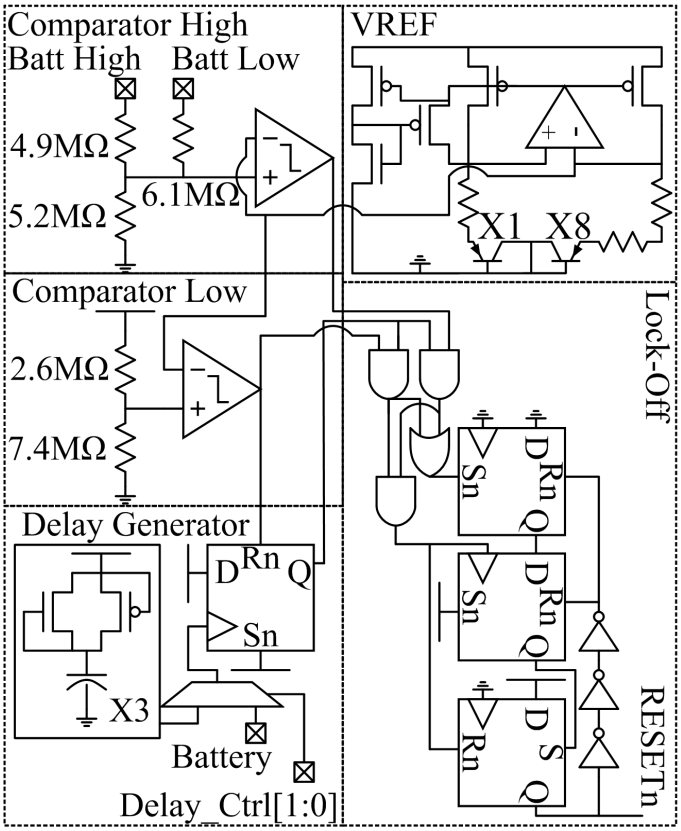

Fig. 12. On-chip 1.7uA POR/BOD with "lock-off" for Carbon-Zinc batteries

Switched Capacitor Networks are the cheaper choice since they only contain capacitors, but they are only capable of having a fixed conversion ratio. In order to overcome this limitation and provide an output range of 3.2V over a wide range of battery input voltages, we created a Step Down SCN which provides fractional voltages in the range of 25%, 33%, 50%, 66%, 75%, 100% of battery voltage (Table 3), and used a Step Up SCN whose output voltage is determined by the output of the Step Down SCN + Battery Voltage.

	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

Table 3. Configurations for Step Down SCN to achieve fractional voltages

The Step Down SCN achieves its fractional voltages of the battery voltage input through configurations of its switches (Fig. 13) (Table 3). This combination of fractional voltages from the Step Down SCN output that is fed into the Step Up SCN that is further boosted by the battery voltage ensures that LIT can maintain a 10% range of 3.2V (Fig. 14). Furthermore, our Hybrid Switch Capacitor Network can be bypassed to output the battery voltage and can be power gated when unused to save energy.

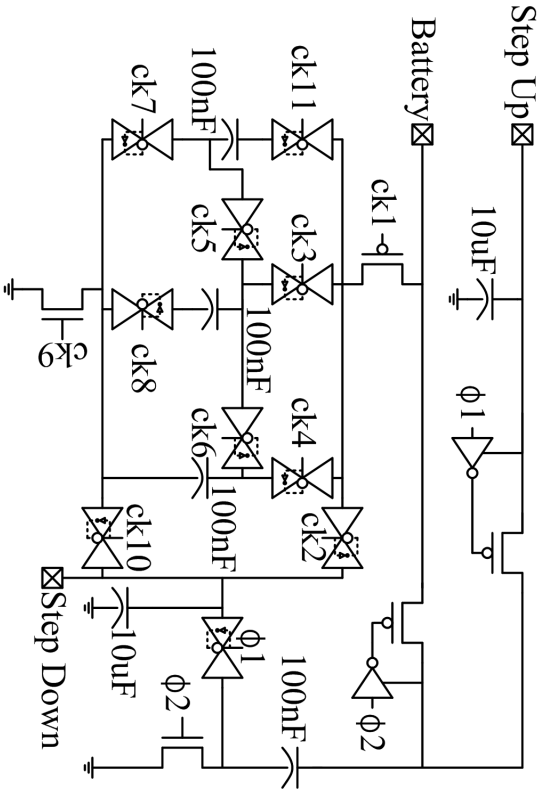

Fig. 13. Variable Step-Up Converter using set-down/-up SCNs.

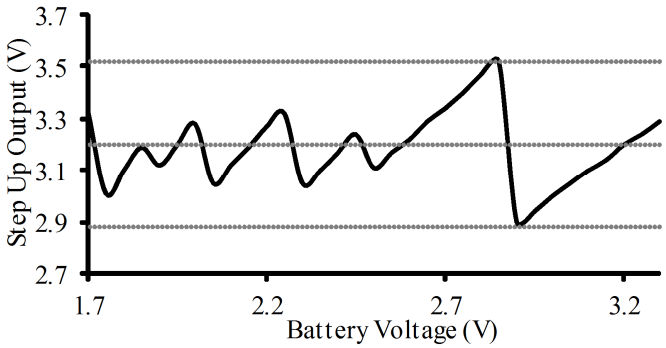

Fig. 14. Measured Output is shown to stay within 10% of 3.2 across 3.3V - 1.7V battery voltage range.

G. Wakeup Interrupt Controller

LIT needs some way of coordinating all its operational modes (Active, Standby, Deep Sleep), managing the power supplies, (LDOs, Direct Battery Connect), 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 mode. Thus, we created the Wakeup Interrupt Controller (WIC) which achieves all the above.

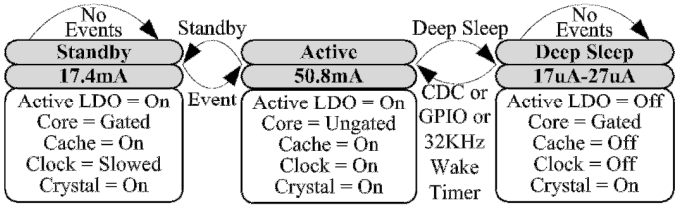

Fig. 15. Wakeup-Interrupt Controller State Diagram

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 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 D. 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 & Dirty LDO and their associated VREFs and IREFs during Deep Sleep mode.

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 Section B.

Lastly, the WIC monitors the events triggers during Deep Sleep mode, such as the touch interface (CDCs), GPIO input, 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.

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. Our 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.

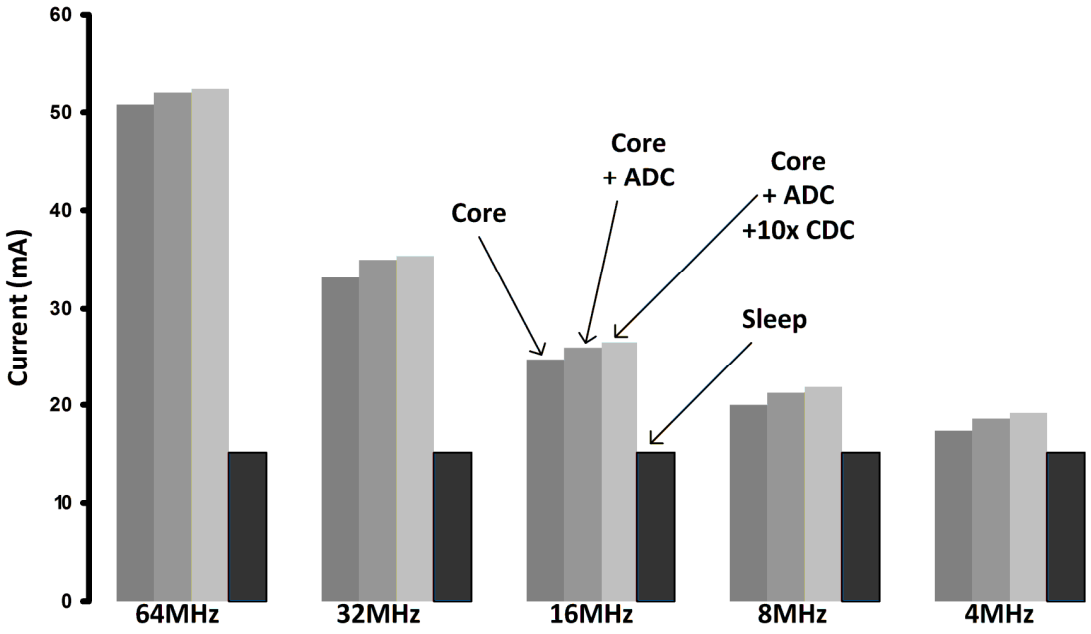

Fig. 16. Silicon Measured Current Draw for Active mode with Operational Blocks and Sleep Mode at various Frequencies

H. Results

LIT's workload can vary greatly depending on its usage, so we can alter its clock frequency, power gate blocks depending on usage, and therefore, 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 (SOLA) 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 (Fig. 16).

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.

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 17uA (32KHz WakeTimer Monitoring) to 27uA (All Event Monitoring) that results in ~2 years of Deep Sleep lifetime.

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

Table 4. Deep Sleep mode current draw with event monitoring options selected

We conclude our results with LIT's silicon measured core frequency vs. Voltage and silicon measured core GOPS/WATT (Fig. 17), Shmoo Plot (Fig. 18), die micrograph (Fig. 19), and design parameters and measured results (Table 5) shown below.

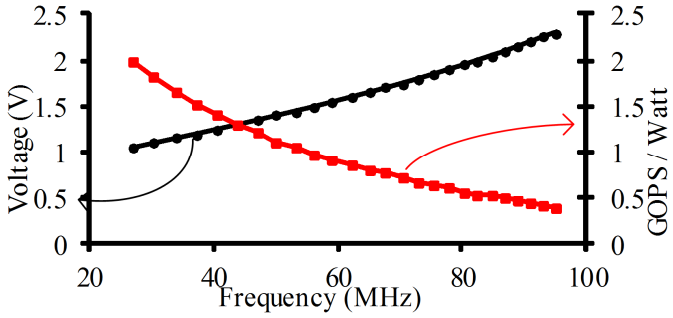

Fig. 17. Silicon Measured Maximum Core Clock Frequency vs. Voltage & Silicon Measured Core GOPS/Watt

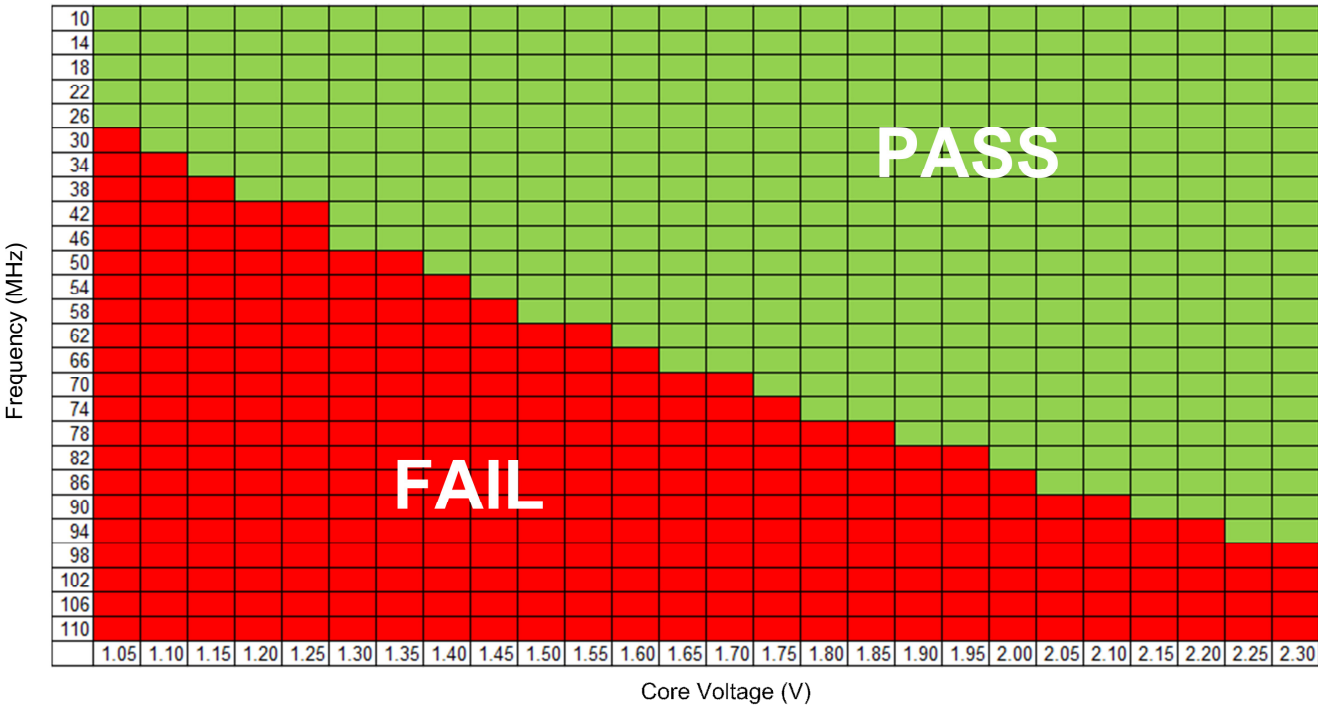

Fig. 18. Shmoo Plot

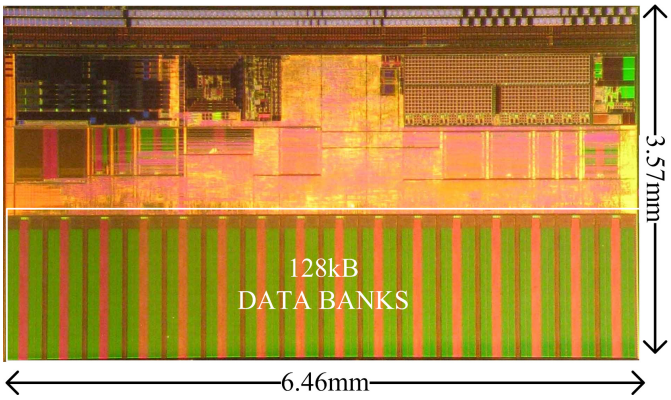

Fig. 19. Die Micrograph

Technology	180nm
Area	23.06mm^2
Performance	64MHz
# Transistors	8Million
# Gates	265K
Cache	128kB True LRU
Cache Area	46%
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	3 sample/s 0.72mW 13ENOB

Near Field Inductive Coil	BER < 10^{-6} 660kbps at 6.5cm
Active Power	91mW
Standby Power	27mW
Deep Sleep Power	30.6uW
Operating Voltage Range	3.3V – 1.7V
Sleep Lifetime (Carbon Zinc)	~2 years

Table 5. Design Parameters and Silicon Measured Results

III. Conclusion

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, Voltage Booster, 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 allows us to provide a low cost, affordable solution for the dissemination of information to illiterate people groups in developing regions.

REFERENCES

[1] One Laptop Per Child, http://one.laptop.org

[2] UNESCO, http://www.unesco.org

[3] Books of Hope, http://booksofhope.com

[4] Global Recording Networks http://globalrecordings.net/topic/saber

[5] Literacy Bridge, http://www.literacybridge.org

[6] C. Schmidt, et al, "Impact of low-cost, on-demand, information access in a remote Ghanaian village," IEEE International Conference on Information and Communication Technologies and Development, 2010.

[7] C.K. Prahalad, "The Fortune at the Bottom of the Pyramid," NJ, Wharton School Publishing, 2006.

[8] ZhiYoong Foo, et al, "A Case for Custom Silicon in Enabling Low-Cost Information Technology for Developing Regions," ACM Symposium on Computing for Development, December 2010.

[9] John F. Dickson, "On-Chip High-Voltage Generation in Integrated Circuits Using an Improved Multiplier Technique," J. Solid-State Circuits, vol. 11, no. 3, pp. 374-378, June 1976.

[10] Ramadass, Y.K.; Chandrakasan, A.P.; , "Voltage Scalable Switched Capacitor DC-DC Converter for Ultra-Low-Power On-Chip Applications," Power Electronics Specialists Conference, 2007. PESC 2007. IEEE , vol., no., pp.2353-2359, 17-21 June 2007.

[11] R. Jacob Baker, "CMOS: Circuit Design, Layout, and Simulation," NY, Wiley-IEEE Press, 2005.

[12] G. Chen, et. al., "A Cubic-Millimeter Energy-Autonomous Wireless Intraocular Pressure Monitor," ISSCC Dig. Tech. Papers, pp. 310-311, 2011.

[13] K. Martin and D. Johns, "Analog Integrated Circuit Design," Wiley, 1997.

[14] Y. Ye, S. Borkar, V. De, "A New Technique for Standby Leakage Reduction in High Performance Circuits," 1998 Symposium on VLSLI Circuits, June 1998, pp.40-41

[15] A. Chandrakasan, W. Bowhill, F. Fox, "Design of High Performance Microprocessor Circuits," NY, Wiley-IEEE Press, 2000

[16] Huy-Binh Le; Xuan-Dien Do; Sang-Gug Lee; 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* , vol.58, no.11, pp.778-782, Nov. 2011

[17] Rincon-Mora, G.A.; Allen, P.E.; , "A low-voltage, low quiescent current, low drop-out regulator," Solid-State Circuits, IEEE Journal of , vol.33, no.1, pp.36-44, Jan 1998.

[18] Kuijk, K.E.; , "A precision reference voltage source," Solid-State Circuits, IEEE Journal of , vol.8, no.3, pp.222-226, June 1973