
WRL Technical Note TN-59



Power Evaluation of Itsy Version 2.4

*Marc A. Viredaz
Deborah A. Wallach*

The Western Research Laboratory (WRL), located in Palo Alto, California, is part of Compaq's Corporate Research group. WRL was founded by Digital Equipment Corporation in 1982. We focus on information technology that is relevant to the technical strategy of the Corporation, and that has the potential to open new business opportunities. Research at WRL includes Internet protocol design and implementation, tools to optimize compiled binary code files, hardware and software mechanisms to support scalable shared memory, graphics VLSI ICs, handheld computing, and more. As part of WRL tradition, we test our ideas by extensive software or hardware prototyping.

We publish the results of our work in a variety of journals, conferences, research reports, and technical notes. This document is a technical note. We use technical notes for rapid distribution of technical material; usually this represents research in progress. Research reports are normally accounts of completed research and may include material from earlier technical notes, conference papers, or magazine articles.

You can retrieve research reports and technical notes via the World Wide Web at:

<http://www.research.compaq.com/wrl/>

You can request printed copies of research reports and technical notes, when available, by mailing your order to us at:

Technical Report Distribution
Compaq Western Research Laboratory
250 University Avenue
Palo Alto, California 94301 U.S.A.

You can also request reports and notes via e-mail. For detailed instructions, put the word "Help" in the Subject line of your message, and mail it to:

WRL-Techreports@pa.dec.com

Power Evaluation of Itsy Version 2.4

Marc A. Viredaz
Deborah A. Wallach

February 2001

Abstract

The *Itsy* pocket computer is a flexible research platform developed at Compaq Computer Corporation's *Western Research Laboratory (WRL)*. The goal of this project is to enable hardware and software research in handheld computing. Several versions of the *Itsy* system were designed and built. This document evaluates the power consumption of *Itsy* version 2.4. A similar study was performed for *Itsy* v2.3. The methodology used for the present study is essentially the same, although it has been improved in several ways in order to reduce measurement errors.

Revision 1.0

Acknowledgements

The Itsy project was started and led by William R. Hamburger. The hardware was designed jointly by William R. Hamburger and one of the authors (Marc A. Viredaz). The Linux operating system was ported to Itsy by Lawrence S. Brakmo, Carl A. Waldspurger, and the second author (Deborah A. Wallach). Joel F. Bartlett and Timothy Mann have also made significant contributions to the operating system and the applications used for the present study. The Itsy project has also benefited from the work of many other researchers at the *Western Research Laboratory (WRL)* and the *System Research Center (SRC)*. Special thanks are due to Wayne Mack, for his help with the experimental setup. The authors would also like to acknowledge the enthusiastic support of their management during this major project.

Contents

Acknowledgements	iv
Contents	v
List of Figures	vi
List of Tables	vi
1 Introduction	1
2 Methodology	1
2.1 Experimental setup	1
2.1.1 Average power	2
2.1.2 Battery lifetime	3
2.2 Errors in the average power measurements	3
2.2.1 Hardware instrumentation errors	4
2.2.2 Averaging error	5
2.2.3 Software instrumentation error	5
2.2.4 Combined average power error	6
2.3 Errors in the battery lifetime measurements	6
2.3.1 Timing error	6
2.3.2 Software initialization error	7
2.3.3 Software instrumentation error	7
2.3.4 Combined battery lifetime error	7
2.4 Experiment-to-experiment differences	7
2.4.1 Average power and average battery lifetime	8
3 Benchmarks	8
4 Results	10
5 Conclusion	12
References	12
A List of symbols	14
B Average power approximation	15
C Multimeter configuration	16
D Complete results	17

List of Figures

1	Experimental setup for Itsy v2.4 power measurements	1
2	Experiment timings	3

List of Tables

1	Benchmark parameters	8
2	Average power and average battery lifetime over all experiments	10
3	Benchmark results : deep sleep	18
4	Benchmark results : sleep	18
5	Benchmark results : sleep, daughter-card (DC)	18
6	Benchmark results : sleep, LCD	18
7	Benchmark results : idle, 59.0 MHz, low voltage (LV)	18
8	Benchmark results : idle, 59.0 MHz	19
9	Benchmark results : idle, 132.7 MHz	19
10	Benchmark results : idle, 206.4 MHz	19
11	Benchmark results : WAV, 59.0 MHz	19
12	Benchmark results : WAV, 206.4 MHz	19
13	Benchmark results : DECTalk, 73.7 MHz, low voltage (LV)	20
14	Benchmark results : DECTalk, 73.7 MHz	20
15	Benchmark results : DECTalk, 206.4 MHz	20
16	Benchmark results : dictation, daughter-card (DC), 206.4 MHz	20
17	Benchmark results : MPEG-1, 206.4 MHz	20

1 Introduction

The aim of the *Itsy* project is to develop flexible research platforms for pocket computing. Several revisions of Itsy version 1 were built between 1997 and 1999 [Vir98a, Vir98b]. Itsy version 2 is the result of a major redesign started in 1999. Of the version 2 Itsy units, Itsy v2.0 refers to the first mother-board *printed-circuit board (PCB)* design, while Itsy v2.1 corresponds to an Itsy v2.0 modified to fix its bugs. A total of five Itsy v2.0/v2.1 units have been built. Similarly, Itsy v2.2 refers to the next revision of the PCB, while Itsy v2.3 refers to a bug-fixed Itsy v2.2. Slightly more than 25 Itsy v2.2/v2.3 units have been built. The final revision is Itsy v2.4. About 125 such systems have been manufactured. The Itsy v2 hardware and software have been described elsewhere [Bar00] and will not be presented in this document.

The purpose of this study is to accurately measure the battery lifetime of Itsy v2.4 under different loads, while continuously monitoring the power consumption. A similar study was performed for Itsy v2.3 [Vir00]. The methodology used for the present study is essentially the same, although it has been improved in several ways in order to reduce measurement errors. Because this report is intended to be a stand-alone document, the complete methodology is described, although much of the information presented here is identical to that found in the previous report [Vir00]. To help readers who are familiar with the Itsy v2.3 study, all differences are explicitly outlined.

2 Methodology

This section introduces the experimental setup used to measure the average power consumption and battery lifetime. An error analysis for this setup is then presented.

2.1 Experimental setup

Figure 1 shows the experimental setup used for power measurements. The Itsy v2 system uses a 4.1 V prismatic Lithium-ion battery of size 48 mm × 30 mm × 6 mm, which is rated at 640 mA · h.

In series with the battery is a precision sense resistor $R = 0.1 \Omega$ [IRC99] used by Itsy's battery monitor [Dal00]. It can also be used by an external instrument, as in this study. Two external HP 34401A multimeters

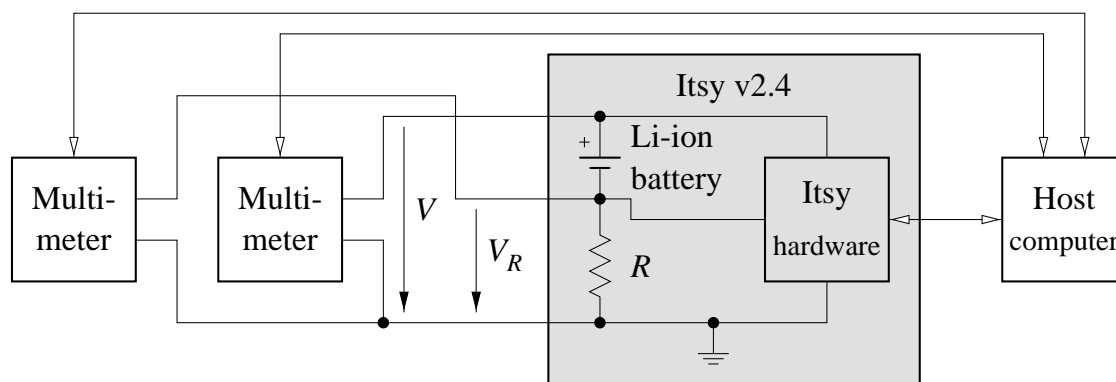


Figure 1: Experimental setup for Itsy v2.4 power measurements.

[HP96a, HP96b] monitor the battery voltage (V) and the voltage drop across the sense resistor (V_R). Both the Itsy and the multimeters are connected to a host computer through RS-232 links.

Both sides of the sense resistor are connected to the rest of the Itsy hardware. Since the connection on the negative pole of the battery is attached to both a high-impedance input and a filter capacitor, its leakage current is negligible (around 10 nA), but the voltage V_R is subject to the battery monitor's RC filter, whose time constant is 4.7 ms.

Although not shown in Figure 1, a power supply is connected to the battery charging circuit through a relay. Another relay is used to reset the Itsy. By controlling both relays, the host computer can independently run a series of experiments, charging the battery between experiments.¹ Since both relays are open during the actual experiments, this additional setup will not be further considered.

2.1.1 Average power

In this document, the term *integration* refers to a single voltage measurement made by a multimeter.² Both instruments are configured to perform N integrations while calculating the average voltage and keeping track of the minimum and maximum values. Once these N integrations have been completed, the values $V_{\text{avg}, a}$, $V_{\text{min}, a}$, $V_{\text{max}, a}$, $V_{R \text{ avg}, a}$, $V_{R \text{ min}, a}$, and $V_{R \text{ max}, a}$ are down-loaded to the host computer. This collection of values is referred to as an *acquisition*. The host computer also records the start and end times $t_{\text{start}, a}$ and $t_{\text{end}, a}$ of each acquisition a .

The integration time has been chosen as $t_{\text{int}} = \frac{0.2}{60 \text{ Hz}} = 3.33 \text{ ms}$. The multimeters are configured to perform a 3.33 ms zero adjustment before each 3.33 ms integration. This zero adjustment — an integration performed with the input disconnected — allows the instrument to compensate for any charge accumulations or other forms of internal drift. Therefore, the approximate duration of an acquisition is:

$$t_{\text{acq}, a} = t_{\text{start}, a+1} - t_{\text{start}, a} \approx 2 \cdot N \cdot t_{\text{int}} \quad (1)$$

Experience showed that long acquisition times³ can result in large averaging errors (see Section 2.2.2). Therefore, an acquisition time of $t_{\text{acq}, a} \approx 30 \text{ s}$ was chosen for all experiments and the number of integrations was empirically selected as $N = 4150$ to achieve this result. The acquisition time predicted by Equation (1) is about 8 % smaller than 30 s, due to up-load and down-load time and the internal overhead of the multimeters.

From Ohm's law, the battery current I is equal to V_R/R , allowing the average, minimum, and maximum currents $I_{\text{avg}, a}$, $I_{\text{min}, a}$, and $I_{\text{max}, a}$ to be calculated. Thus, for each acquisition a , an approximation $\tilde{P}_{\text{avg}, a}$ of the average power $P_{\text{avg}, a}$ is:

$$\tilde{P}_{\text{avg}, a} = V_{\text{avg}, a} \cdot I_{\text{avg}, a} \quad (2)$$

while upper and lower bounds on the power are:

$$P_{\text{low bnd}, a} = V_{\text{min}, a} \cdot I_{\text{min}, a} \quad (3)$$

$$P_{\text{up bnd}, a} = V_{\text{max}, a} \cdot I_{\text{max}, a} \quad (4)$$

¹ For the Itsy v2.3 power evaluation, the sleep-mode experiments were started manually. For the present study, they have been modified to start from the Linux operating system. Therefore, the complete suite of benchmarks can be run automatically.

² The multimeter's electronics integrates the voltage over a short period of time, thus resulting in an average over this period.

³ For the Itsy v2.3 power evaluation, 5-minute acquisitions were used for sleep-mode and idle-mode experiments.

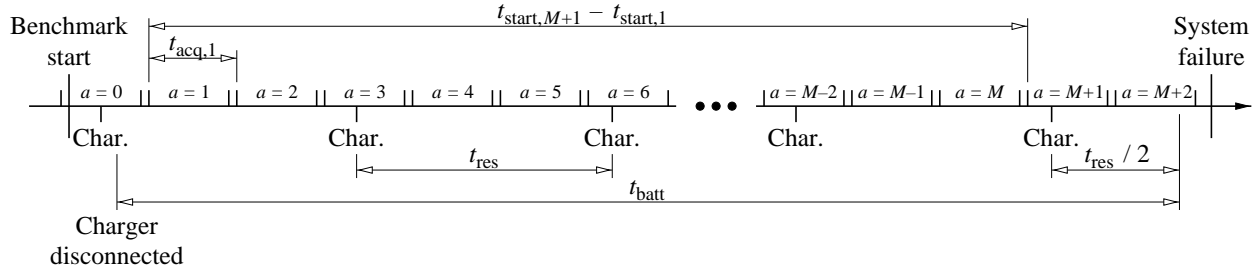


Figure 2: Experiment timings.

Finally, the average power P_{avg} over the battery lifetime is the weighted mean of all the acquisitions:

$$P_{\text{avg}} = \frac{\sum_{a=1}^M t_{\text{acq},a} \cdot P_{\text{avg},a}}{\sum_{a=1}^M t_{\text{acq},a}} = \frac{\sum_{a=1}^M t_{\text{acq},a} \cdot P_{\text{avg},a}}{t_{\text{start},M+1} - t_{\text{start},1}} \approx \frac{\sum_{a=1}^M t_{\text{acq},a} \cdot \tilde{P}_{\text{avg},a}}{t_{\text{start},M+1} - t_{\text{start},1}} \quad (5)$$

where M is the number of “meaningful” acquisitions. Only acquisitions which start after the charger has been disconnected are considered. By convention, the first such acquisition is referred to as $a = 1$. Also, since the battery lifetime does usually not correspond to an integer number of acquisitions, only acquisitions that are guaranteed to be completely within the battery lifetime are kept. This is illustrated in Figure 2.

2.1.2 Battery lifetime

At regular intervals t_{res} , the Itsy sends a string of two characters (i.e., a character followed by a new line) to the host computer. In contrast to the Itsy v2.3 power evaluation, a benchmark is started while the Itsy is still connected to the charging power supply. Once the benchmark is running, as indicated by the first string of characters received, the charger is disconnected. The experiment is considered finished when the characters are no longer received by the host (i.e., the Itsy has entered sleep mode due to a battery fault).

The resolution interval t_{res} represents a tradeoff between the overhead to send characters (see Section 2.2.3 and 2.3.3) and the resolution to which the battery lifetime is measured (see Section 2.3.1). It is chosen separately for each experiment. In order to make the associated error symmetric, the battery lifetime t_{batt} is defined to start when the charger is disconnected and to end half of the resolution time (i.e., $t_{\text{res}}/2$) after the last string of characters is received, as shown in Figure 2.

2.2 Errors in the average power measurements

The goal of this section is to find an upper bound on the error $\varepsilon_{P_{\text{avg}}}$ on the average power P_{avg} . In the following analysis, all errors are expressed as relative errors, represented by the symbol ε . When absolute errors are needed, the relative error is simply multiplied by its associated value. This section considers only possible errors made during a particular experiment on a particular unit. Measurement-to-measurement and system-to-system differences are discussed in Section 2.4. All potentially important sources of errors are analyzed and the significant terms are kept in the final combined error $\varepsilon_{P_{\text{avg}}}$.

2.2.1 Hardware instrumentation errors

The tolerance of the sense resistor R is $\varepsilon_R = 1\%$ [IRC99]. The PCB has been routed such that the test points are connected directly to the resistor's pad, opposite the main current path, to avoid measuring any voltage drop due to the resistivity of the main-current traces.

Although the multimeters are set on auto-range (see Appendix C), the 10 V range is always used to measure the battery voltage V and the 100 mV range is always used to measure the sense resistor voltage V_R . According to the user's manual [HP96a], the relative errors on these measurements are:

$$\varepsilon_V = \frac{50 \mu\text{V}}{V} + 0.0035\% \quad (6)$$

$$\varepsilon_{V_R} = \frac{3.5 \mu\text{V}}{V_R} + 0.0050\% \quad (7)$$

Because the absolute error of an average is the average of the individual absolute errors, it can easily be seen that the same formulae also apply to the relative errors for the average voltages:

$$\begin{aligned} \varepsilon_{V_{\text{avg}, a}} &= \frac{\frac{1}{N} \cdot \sum_{i=1}^N \left(\frac{50 \mu\text{V}}{V_i} + 0.0035\% \right) \cdot V_i}{V_{\text{avg}, a}} = \frac{50 \mu\text{V} + 0.0035\% \cdot \frac{1}{N} \cdot \sum_{i=1}^N V_i}{V_{\text{avg}, a}} \\ &= \frac{50 \mu\text{V}}{V_{\text{avg}, a}} + 0.0035\% \end{aligned} \quad (8)$$

similarly:

$$\varepsilon_{V_{R \text{ avg}, a}} = \frac{3.5 \mu\text{V}}{V_{R \text{ avg}, a}} + 0.0050\% \quad (9)$$

In practice, the battery-voltage error $\varepsilon_{V_{\text{avg}, a}}$ is negligible, since it is always smaller than 0.006%. On the other hand, the resistor-voltage error $\varepsilon_{V_{R \text{ avg}, a}}$ is important at low power, particularly in deep-sleep mode, where it is around 3%.⁴

It can be seen that the value of the sense resistor is a compromise between accuracy at low currents and wasted power at high currents. It would have been possible to change the value of the sense resistor depending on the benchmark (i.e., use a higher resistance for low-power benchmarks). However, for the sake of simplicity, it was considered preferable to keep the resistor provided by the Itsy hardware.

Current leakages through the multimeters are negligible, since the input impedance of the multimeters, as configured (see Appendix C), is greater than 10 G Ω for the 100 mV to 10 V ranges [HP96a].

Another source of error is due to the periods of time when the Itsy is not monitored, as explained in Section 2.1.1 (zero adjustment and 8% overhead). Although this error is not quantifiable—at least not without adding another source of error (i.e., internal drifts) or using a more sophisticated test system—it is assumed to be small. This can be justified by the fact that all the benchmarks used are of a constant nature (e.g., sleep mode, idle mode) or of a cyclic nature (e.g., computation loops) with each experiment consisting of a large number of cycles. It is extremely unlikely that the benchmark could be in synchronization with

⁴ This error is higher than the corresponding one for Itsy v2.3, because the sleep-mode power consumption has been reduced.

the measurements. Moreover, even if this were the case (i.e., the benchmark's period was equal to some multiple of the time to perform a zero adjustment and an integration), the fact that the measurement results are up-loaded to the host computer after each acquisition should introduce enough randomness to prevent any constant phase relationship. Therefore, the long-term average should be representative of the workload.

2.2.2 Averaging error

As mentioned in Section 2.1.1, only an approximation $\tilde{P}_{\text{avg}, a}$ of the average power $P_{\text{avg}, a}$ can be calculated for each acquisition a . However, it is possible to find an upper bound on the difference between these two figures. It can be shown (see Appendix B) that:

$$\tilde{P}_{\text{avg}, a} = P_{\text{avg}, a} + \frac{\sum_{i=1}^{N-1} \left(\sum_{j=i+1}^N (V_i - V_j) \cdot (I_j - I_i) \right)}{N^2} \quad (10)$$

Since the double summation in Equation (10) has $N \cdot (N - 1)/2$ terms, this difference can be bounded by:

$$\left| \tilde{P}_{\text{avg}, a} - P_{\text{avg}, a} \right| \leq \frac{N - 1}{2 \cdot N} \cdot (V_{\text{max}, a} - V_{\text{min}, a}) \cdot (I_{\text{max}, a} - I_{\text{min}, a}) \quad (11)$$

Therefore, the relative error of using the approximation $\tilde{P}_{\text{avg}, a}$ instead of the average power $P_{\text{avg}, a}$ is at most:

$$\varepsilon_{\text{avg}, a} = \frac{1}{\tilde{P}_{\text{avg}, a}} \cdot \frac{N - 1}{2 \cdot N} \cdot (V_{\text{max}, a} - V_{\text{min}, a}) \cdot (I_{\text{max}, a} - I_{\text{min}, a}) \quad (12)$$

2.2.3 Software instrumentation error

As mentioned in Section 2.1.2, a string of two characters is sent at regular intervals t_{res} . Experiments have shown that the host computer does not always receive the first characters correctly if the bit rate is too high. Therefore, the highest speed at which no character is ever lost is used: 19200 bit/s (8-bit data, no parity, 1 start bit, 1 stop bit). Hence, the time to send the string of characters is: $\frac{2 \cdot (8+2)}{19200} = 1.04$ ms.

In the sleep-mode experiments, the processor must exit sleep mode, re-initialize part of the hardware, send the string of characters, and re-enter sleep mode. According to the StrongARM SA-1100 manual [Int99] it takes approximately 160 ms to exit sleep mode (10 ms for the power to ramp up and 150 ms for the oscillator to stabilize) and 150 μs to enter it. Thus, approximately 160 ms is required for the complete operation. With respect to the Itsy v2.3 power evaluation, this procedure has been optimized by initializing only the minimum amount of hardware. In particular, the DRAM is left unpowered (deep sleep) or in self-refresh (sleep). A few differential measurements on the three systems used for this study showed that the additional energy to execute this operation is between 13 mJ and 16.5 mJ, for all sleep-mode benchmarks.⁵ This value has been rounded to $E_{\text{char}} = 20$ mJ to account for possible variations

⁵ This value was calculated by comparing the average power required to send a string of characters every second to the average power required to send characters much less frequently.

For idle-mode and run-mode experiments, similar differential measurements under the Linux operating system have shown that the cost of sending this string of characters is 40–65 mJ, rounded to $E_{\text{char}} = 80$ mJ. The relative error on the average power P_{avg} due to this instrumentation is equal to:

$$\varepsilon_{\text{char}, P_{\text{avg}}} = \frac{1}{P_{\text{avg}}} \cdot \frac{\left(\left\lfloor \frac{t_{\text{batt}}}{t_{\text{res}}} \right\rfloor - 1 \right) \cdot E_{\text{char}}}{t_{\text{start}, M+1} - t_{\text{start}, 1}} \quad (13)$$

It should be stressed that this overhead is always conservative, since it increases the power consumed during the experiment.

2.2.4 Combined average power error

The final step of the error analysis is to combine all sources of errors. Ignoring second-order effects, the relative error of a product is the sum of the individual relative errors. Likewise, when the error on a given value comes from several sources, the relative errors should be added. Therefore, based on the sense resistor's tolerance and on Equations (8), (9), and (12), the error on the average power $\tilde{P}_{\text{avg}, a}$ for each acquisition a is:

$$\varepsilon_{\tilde{P}_{\text{avg}, a}} \approx \varepsilon_R + \varepsilon_{V_{\text{avg}, a}} + \varepsilon_{V_{R_{\text{avg}, a}}} + \varepsilon_{\text{avg}, a} \quad (14)$$

The next step is to find the error on the weighted average of all $\tilde{P}_{\text{avg}, a}$ terms (note that the absolute error of a weighted average is the weighted average of the individual absolute errors). Finally, the result of Equation (13) must be added to calculate the total error on the average power P_{avg} :

$$\varepsilon_{P_{\text{avg}}} \approx \frac{1}{P_{\text{avg}}} \cdot \frac{\sum_{a=1}^M t_{\text{acq}, a} \cdot \varepsilon_{\tilde{P}_{\text{avg}, a}} \cdot \tilde{P}_{\text{avg}, a}}{t_{\text{start}, M+1} - t_{\text{start}, 1}} + \varepsilon_{\text{char}, P_{\text{avg}}} \quad (15)$$

2.3 Errors in the battery lifetime measurements

This section presents a similar error analysis, aimed at finding an upper bound on the error $\varepsilon_{t_{\text{batt}}}$ on the battery lifetime t_{batt} .

2.3.1 Timing error

As mentioned in Section 2.1.2, the resolution at which the battery lifetime t_{batt} is measured is the interval t_{res} between strings of characters sent to the host computer. The corresponding relative error is therefore:

$$\varepsilon_{\text{res}} = \frac{1}{t_{\text{batt}}} \cdot \frac{t_{\text{res}}}{2} \quad (16)$$

Given that the resolution interval t_{res} is between 90 s and 1 hour, the precision of the host computer's clock is negligible.

2.3.2 Software initialization error

In this study, the software initialization error that affected the Itsy v2.3 power evaluation has been eliminated ($\varepsilon_{\text{init}} = 0$) by keeping the battery charging until after a benchmark has started. A very small error could be introduced if the power required during the initialization is higher than the maximum power that the charger circuit can provide and, therefore, some power is drained from the battery. However, this error can safely be ignored since it is much smaller than other terms and it is conservative.

2.3.3 Software instrumentation error

The fact that a string of characters is sent at regular intervals to the host computer (see Section 2.1.2) also introduces an error on the battery lifetime t_{batt} . This error is derived in a similar way as the corresponding error on the average power (see Section 2.2.3). In the case of the battery lifetime, the duration to be considered is t_{batt} rather than $t_{\text{start}, M+1} - t_{\text{start}, 1}$. Therefore, one more character string is sent, as shown in Figure 2. Therefore, the relative error on the battery lifetime t_{batt} due to this instrumentation is equal to:

$$\varepsilon_{\text{char}, t_{\text{batt}}} = \frac{1}{t_{\text{batt}}} \cdot \frac{\left\lfloor \frac{t_{\text{batt}}}{t_{\text{res}}} \right\rfloor \cdot E_{\text{char}}}{P_{\text{avg}}} \quad (17)$$

Note that the first character string is sent during the initialization phase and is not accounted for. Mathematically, Equation (17) is correct only when the interval between this character string and the starting point of t_{batt} is smaller than $t_{\text{res}}/2$, which is always the case in this study.

2.3.4 Combined battery lifetime error

The combined error on the battery lifetime t_{batt} is simply the sum of Equations (16) and (17):

$$\varepsilon_{t_{\text{batt}}} \approx \varepsilon_{\text{res}} + \varepsilon_{\text{char}, t_{\text{batt}}} \quad (18)$$

2.4 Experiment-to-experiment differences

The analysis presented in Sections 2.2 and 2.3 evaluates the error on a single experiment. However, the result of similar experiments will differ among different units and even among measurements using the same unit. These differences could be evaluated by statistical analysis. However, because of the length of time required to execute the complete benchmark suite (about two months, including battery charging), only three sets of experiments (i.e., as many as possible given the available test equipment) have been run and a formal statistical analysis has not been conducted.

Regarding measurement-to-measurement differences (on the same unit), one important factor is battery aging. In order to minimize this influence, all batteries used in this study had never been used before.⁶ Therefore, the results presented here are relevant for new batteries. Battery lifetime is expected to decrease as the battery is being used. Based on the manufacturer's data, prismatic Lithium-ion cells typically retain 70% – 90% of their initial capacity after 500 low-current charge/discharge cycles.

⁶ Specifically, the batteries had been stored as received (i.e., about half charged) at approximately 4°C and were about 22 months past their date of manufacture.

Experiment	Freq. [MHz]	N	$t_{\text{acq}, a}$ [s]	E_{char} [mJ]	t_{res} [s]
Deep sleep	–	4150	30	20	3600
Sleep	–	4150	30	20	3600
Sleep, DC	–	4150	30	20	3600
Sleep, LCD	–	4150	30	20	3600
Idle, 59.0 MHz, LV	59.0	4150	30	80	600
Idle, 59.0 MHz	59.0	4150	30	80	600
Idle, 132.7 MHz	132.7	4150	30	80	600
Idle, 206.4 MHz	206.4	4150	30	80	600
WAV, 59.0 MHz	59.0	4150	30	80	252
WAV, 206.4 MHz	206.4	4150	30	80	252
DECTalk, 73.7 MHz, LV	73.7	4150	30	80	195
DECTalk, 73.7 MHz	73.7	4150	30	80	195
DECTalk, 206.4 MHz	206.4	4150	30	80	195
Dictation, DC, 206.4 MHz	206.4	4150	30	80	102
MPEG-1, 206.4 MHz	206.4	4150	30	80	90

Table 1: Benchmark parameters.

2.4.1 Average power and average battery lifetime

Section 4 presents the reported average power $P_{\text{avg}}^{\text{base}}$ and average battery lifetime $t_{\text{batt}}^{\text{base}}$, over all experiments. These are simply the arithmetic means of the average power P_{avg} and battery lifetime $t_{\text{batt}}^{\text{base}}$. Their respective relative errors have been conservatively calculated to include all of the individual error ranges:

$$\varepsilon_{P_{\text{avg}}^{\text{base}}} = \frac{1}{P_{\text{avg}}^{\text{base}}} \cdot \max \left(P_{\text{avg}}^{\text{base}} - \min_x \left((1 - \varepsilon_{P_{\text{avg}}^x}) \cdot P_{\text{avg}}^x \right), \max_x \left((1 + \varepsilon_{P_{\text{avg}}^x}) \cdot P_{\text{avg}}^x \right) - P_{\text{avg}}^{\text{base}} \right) \quad (19)$$

$$\varepsilon_{t_{\text{batt}}^{\text{base}}} = \frac{1}{t_{\text{batt}}^{\text{base}}} \cdot \max \left(t_{\text{batt}}^{\text{base}} - \min_x \left((1 - \varepsilon_{t_{\text{batt}}^x}) \cdot t_{\text{batt}}^x \right), \max_x \left((1 + \varepsilon_{t_{\text{batt}}^x}) \cdot t_{\text{batt}}^x \right) - t_{\text{batt}}^{\text{base}} \right) \quad (20)$$

where x is the experiment index.

3 Benchmarks

This section describes the workload used for the power experiments. All benchmarks were executed on an Itsy without a daughter-card, except when otherwise specified. All idle-mode and run-mode experiments were run under the Linux operating system, with the serial interfaces (USB, RS-232, infrared, etc.) disabled, the audio disabled when unused, and the other I/O devices (LCD, touch screen, etc.) left enabled. This configuration simulates a system as typically used when not connected to a host computer. In all experiments, the RS-232 interface is only enabled during the short period of time when the processor sends a string of characters (see Section 2.1.2). Table 1 shows the parameters used for all benchmarks.

The following benchmarks were used:

Deep sleep

In this benchmark, the processor is in sleep mode with the oscillator disabled [Int99], the DRAM is unpowered, the flash memory is in standby mode, and all I/O devices (serial interfaces, LCD, audio, etc.) are disabled. When the processor exits sleep mode to send a string of characters, it is clocked at 206.4 MHz.

Sleep

In this benchmark, the processor is in sleep mode with the oscillator disabled [Int99], the DRAM is in self-refresh mode, the flash memory is in standby mode, and all I/O devices (serial interfaces, LCD, audio, etc.) are disabled. When the processor exits sleep mode to send a string of characters, it is clocked at 206.4 MHz. This benchmark is executed both on an Itsy alone and on an Itsy with a memory-extension *daughter-card (DC)* providing an additional 32 Mbyte of flash memory and 32 Mbyte of DRAM.

Sleep, LCD

In this benchmark, the processor is in sleep mode with the oscillator disabled [Int99], the DRAM is in self-refresh mode, the flash memory is in standby mode, and all I/O devices (serial interfaces, audio, etc.) are disabled, except the LCD, which is configured to display a static monochrome image. When the processor exits sleep mode to send a string of characters, it is clocked at 206.4 MHz. This benchmark was not included in the Itsy v2.3 power evaluation.

Idle

In this benchmark, the Itsy is running the Linux operating system without a workload. Therefore, the processor is mostly in idle mode. This benchmark is run at 59.0 MHz, 132.7 MHz, and 206.4 MHz. At 59.0 MHz, it is executed both at the standard core voltage of 1.5 V and at a *low voltage (LV)* of 1.23 V.

WAV

In this benchmark, the Itsy is playing an audio file in WAV format, using a simple program specially written for Itsy. The file played is the audio track to the video file used in the MPEG experiment. The built-in speaker is used and the audio volume is set to the maximum possible (the speaker dissipates about 135 mW). This benchmark is run at 59.0 MHz and 206.4 MHz.

DECTalk

In this benchmark, the Itsy is generating an audio stream from a text file (i.e., text-to-speech), using the DECTalk program. During this experiment, the built-in speaker only dissipates about 25–35 mW. This benchmark is run at 73.7 MHz and 206.4 MHz. At 73.7 MHz, it is executed both at the standard core voltage of 1.5 V and at a *low voltage (LV)* of 1.23 V.

Experiment	Num. exp.	Idle time	$P_{\text{avg}}^{\text{base}}$ [mW]	$\varepsilon_{P_{\text{avg}}^{\text{base}}}$	$t_{\text{batt}}^{\text{base}}$ [h] ([days])	$\varepsilon_{t_{\text{batt}}^{\text{base}}}$
Deep sleep	3	–	4.58	9.5 %	500.0 (20.8)	4.3 %
Sleep	3	–	7.40	5.8 %	308.5 (12.9)	3.5 %
Sleep, DC	3	–	10.6	3.1 %	215.0 (9.0)	2.0 %
Sleep, LCD	3	–	26.2	2.6 %	87.0 (3.6)	2.6 %
Idle, 59.0 MHz, LV	3	95 %	55.4	4.7 %	40.6	4.9 %
Idle, 59.0 MHz	3	95 %	69.5	6.6 %	32.3	6.8 %
Idle, 132.7 MHz	3	95 %	83.0	6.1 %	27.0	6.3 %
Idle, 206.4 MHz	3	95 %	101	5.7 %	22.0	6.1 %
WAV, 59.0 MHz	3	83 %	278	3.5 %	7.75	4.2 %
WAV, 206.4 MHz	3	93 %	310	4.1 %	6.88	4.4 %
DECTalk, 73.7 MHz, LV	3	<1 %	352	2.4 %	6.11	3.0 %
DECTalk, 73.7 MHz	3	<1 %	397	2.7 %	5.35	3.4 %
DECTalk, 206.4 MHz	3	53 %	401	3.9 %	5.29	3.7 %
Dictation, DC, 206.4 MHz	3	<1 %	757	3.8 %	2.67	3.5 %
MPEG-1, 206.4 MHz	3	16 %	821	4.8 %	2.42	3.8 %

Table 2: Average power and average battery lifetime over all experiments.

Dictation

In this benchmark, the Itsy is performing continuous-speech recognition (i.e., the translation of dictation) on a prerecorded input file, using the Dragon NaturallySpeaking engine. The recognition engine runs about 2.4 times slower than real time. This benchmark is run at 206.4 MHz. This benchmark is executed on an Itsy with a memory-extension *daughter-card* (DC), providing an additional 32 Mbyte of flash memory and 32 Mbyte of DRAM.

MPEG-1

In this benchmark, the Itsy is playing a video file in MPEG-1 format, using a modified version of the open-source *easympeg* program and *mpeg-lib* libraries. The audio is produced by playing an audio WAV file in a synchronized fashion, using the same program the WAV experiments use. The built-in speaker is used and the audio volume is set to the maximum (the speaker dissipates about 130 mW). This benchmark is run at 206.4 MHz.

4 Results

This section presents a summary and analysis of the results. The detailed average power P_{avg} and battery lifetime t_{batt} numbers for all experiments are reported in Appendix D.

In this document, all physical quantities (power, time, etc.) were rounded to the closest least-significant shown digit. On the other hand, errors were rounded to the next highest such digit. Table 2 presents the average power $P_{\text{avg}}^{\text{base}}$ and the average battery lifetime $t_{\text{batt}}^{\text{base}}$ over all experiments. The third column of Table 2 shows the idle time, as reported by the Linux operating system, for the corresponding benchmark.

This number is not very precisely measured, but is presented to give a better sense of the corresponding workload.

As expected, putting the system into sleep mode dissipates the least power. Even with the significant delay required to exit sleep mode, it is clearly advantageous to put the system in sleep mode whenever possible. When the system is not in sleep mode, the clock frequency and core voltage greatly affect the power consumption, even when the system is mostly idle (as in the idle experiments).

The power of playing a WAV file in real time decreases significantly with the clock frequency. In this case, there is only a small amount of computation to be performed per time unit and the system is mostly idle. Therefore the gain from reduced clock frequency is almost the same as in idle mode. In fact about 2/3 of the power difference between the WAV benchmark and idle mode is used by the speaker. On the other hand, the DECTalk benchmark, which is also a real-time application, does not benefit from reducing the clock frequency. In this case the reduced power while computing at a low frequency is, at least partially, offset by the fact that there is no longer any idle time. Other factors, such as the energy cost of memory accesses, also play an important role. Ideally, this cost should be constant; however, since memory-access timings are implemented as a multiple of the clock period and conservatively rounded up, memory accesses are less efficient (i.e., they last longer and, hence, use more energy) when the Itsy is running at a low frequency.

The Dictation experiment, a continuous speech recognizer, really stretches the Itsy's capabilities. Because the recognizer runs about 2.4 times slower than real time, and the battery lasts only 2.7 h, an Itsy can only support slightly more than an hour of actual dictation. The last benchmark, 16 frame-per-second MPEG-1 video with sound, is the most power hungry. Even at this high rate, the battery lasts for 2.4 h.

As expected, the power consumption during sleep mode has significantly improved on Itsy v2.4 with respect to Itsy v2.3 [Vir00], while the power consumption during idle and run modes remained essentially the same. This is due to the fact that the two-axis accelerometer [Bar00] is powered down during sleep mode on Itsy v2.4, while it remains powered on Itsy v2.3. The sleep-mode average power figures $P_{\text{avg}}^{\text{base}}$ for Itsy v2.4 are smaller by about 1 mW, resulting in a power decrease of 8 % to 15 %, depending on the type of sleep mode. Similarly, the battery lifetime increased by 6 % to 16 %. On the other hand, the idle-mode and run-mode average power figures $P_{\text{avg}}^{\text{base}}$ are within ± 1 % for both versions, except for the dictation benchmark where the discrepancy reaches 2.6 %. Nonetheless, this difference is smaller than the corresponding measurement errors (see Table 16) and might, therefore, not be significant. However, the battery lifetime is systematically worse by 2 % to 6 % on Itsy v2.4 for the idle-mode and run-mode experiments. This is probably due to the fact that on most of the Itsy v2.3 systems, the run-mode and idle-mode experiments were performed before the sleep-mode experiments, while on Itsy v2.4 the sleep-mode benchmarks without daughter-card were always executed first (see Appendix D).

An interesting phenomenon happened during the deep-sleep benchmark on Itsy 34: the minimum power $P_{\text{low bnd}}$ is negative (see Appendix D). This occurred once, towards the end of the battery lifetime, when the main power supply switched from the 3.05 V regulating mode to the input-tracking mode. At this point, a brief surge of current must have flown backwards towards the battery. The same phenomenon either did not happen for the other deep-sleep benchmarks or happened while the multimeter was not recording (during a zero adjustment or an overhead period).

The global errors $\varepsilon_{P_{\text{avg}}}^{\text{base}}$ and $\varepsilon_{t_{\text{batt}}}^{\text{base}}$ reported in Table 2 are significantly higher than the individual errors $\varepsilon_{P_{\text{avg}}}$ and $\varepsilon_{t_{\text{batt}}}$ presented in Appendix D. This is due to the way the global errors are calculated (see Equations (19) and (20)) and reflects system-to-system differences.

It can be seen in Appendix D that the large averaging errors (see Section 2.2.2) and the resulting large

average power errors $\varepsilon_{P_{\text{avg}}}$ that affected the sleep-mode and idle-mode experiments on Itsy v2.3 were greatly improved (see Section 2.1.1).

The average power error $\varepsilon_{P_{\text{avg}}}$ for deep sleep is still large (about 4%). This is mainly due to the multi-meter error $\varepsilon_{V_{R_{\text{avg}}, a}}$, given by Equation (9). This error could have been reduced by changing the value R of the sense resistor. However, the standard resistor was kept because of the advantage of using an unmodified Itsy combined with the fact that deep sleep is not considered a very important mode.⁷

5 Conclusion

To fulfill the Itsy project goals, sufficient processing power and memory capacity were needed to run next-generation applications and user interfaces, as well as sufficient battery lifetime to allow running realistic user interface experiments. This study shows that the power requirements of the Itsy version 2.4 meets these goals for a wide variety of scenarios.

The results presented in Section 4 show that the relationship between clock frequency and power is not intuitive and must be measured. This in turn implies that a good handheld system should have a way to assess its own power consumption, as the Itsy does. It is probably not necessary to have precise measurements. In fact, absolute measurements might not even be necessary: relative measurements might suffice. However, it is important not to forget that the goal of managing power on a battery-powered system is to reduce *energy consumption* not power consumption. For all real-time applications (like most of the benchmarks used in this study), both are equivalent, since their duration is by definition constant. This is not the case for other types of applications.

References

- [Bar00] Joel F. Bartlett, Lawrence S. Brakmo, Keith I. Farkas, William R. Hamburgren, Timothy Mann, Marc A. Viredaz, Carl A. Waldspurger, and Deborah A. Wallach. *The Itsy Pocket Computer*. Research report 2000/6, WRL, Compaq, Palo Alto, CA (USA), October 2000.
- [Dal00] Dallas Semiconductor, Dallas, TX (USA). *DS2437: Smart Battery Monitor*, July 2000.
- [HP96a] Hewlett-Packard. *HP 34401A Multimeter: User's Guide*, 4th edition, February 1996. Manual no. 34401-90004.
- [HP96b] Hewlett-Packard. *HP 34401A Multimeter: Service Guide*, 4th edition, February 1996. Manual no. 34401-90013.
- [Int99] Intel. *Intel® StrongARM® SA-1100 Microprocessor: Developer's Manual*, August 1999. Document no. 278088-004.
- [IRC99] IRC, Corpus Christi, TX (USA). *Low Value Flat Chip Resistor: LRC/LRF Series*, May 1999.
- [Vir98a] Marc A. Viredaz. *The Itsy Pocket Computer Version 1.5: User's Manual*. Technical note TN-54, WRL, Compaq, Palo Alto, CA (USA), July 1998. Revision 1.0.

⁷ Only seldom it is desirable to keep the system powered yet not preserve the DRAM contents.

- [Vir98b] Marc A. Viredaz. *The Itsy Pocket Computer Version 1.5: Hardware Description*. WRL, Compaq, Palo Alto, CA (USA), July 1998. Revision 1.0.
- [Vir00] Marc A. Viredaz and Deborah A. Wallach. *Power Evaluation of Itsy Version 2.3*. Technical note TN-57, WRL, Compaq, Palo Alto, CA (USA), October 2000. Revision 1.0.

A List of symbols

a	Acquisition index.
E_{char}	Additional energy required to send a string of two characters.
I	Battery current.
$I_{\text{avg}, a}$	Average battery current during the acquisition a .
I_i, I_j	Battery current during the multimeter integration i, j .
$I_{\text{max}, a}$	Maximum battery current during the acquisition a .
$I_{\text{min}, a}$	Minimum battery current during the acquisition a .
M	Number of acquisitions used to compute P_{avg} (i.e., “meaningful” acquisitions).
N	Number of multimeter integrations per acquisition.
$P_{\text{avg}}, P_{\text{avg}}^x$	Average battery power during an experiment, and during the experiment x .
$P_{\text{avg}, a}$	Average battery power during the acquisition a .
$\tilde{P}_{\text{avg}, a}$	Approximation of the average battery power during the acquisition a .
$P_{\text{avg}}^{\text{base}}$	Average battery power over all experiments.
$P_{\text{low bnd}}$	Lower bound on the battery power during an experiment.
$P_{\text{low bnd}, a}$	Lower bound on the battery power during the acquisition a .
$P_{\text{up bnd}}$	Upper bound on the battery power during an experiment.
$P_{\text{up bnd}, a}$	Upper bound on the battery power during the acquisition a .
R	Resistance of the sense resistor.
$t_{\text{acq}, a}$	Duration of the acquisition a (until the start of the acquisition $a + 1$).
$t_{\text{batt}}, t_{\text{batt}}^x$	Battery lifetime for an experiment, and for the experiment x .
$t_{\text{batt}}^{\text{base}}$	Average battery lifetime over all experiments.
$t_{\text{end}, a}$	End time of the acquisition a .
t_{init}	Experiment initialization time.
t_{int}	Duration of a multimeter integration.
t_{res}	Resolution interval used to measure t_{batt} .
$t_{\text{start}, a}$	Start time of the acquisition a .
V	Battery voltage.
$V_{\text{avg}, a}$	Average battery voltage during the acquisition a .
V_i, V_j	Battery voltage during the multimeter integration i, j .
$V_{\text{max}, a}$	Maximum battery voltage during the acquisition a .
$V_{\text{min}, a}$	Minimum battery voltage during the acquisition a .
V_R	Voltage drop across the sense resistor.
$V_{R \text{ avg}, a}$	Average voltage drop across the sense resistor during the acquisition a .
$V_{R, i}, V_{R, j}$	Voltage drop across the sense resistor during the multimeter integration i, j .
$V_{R \text{ max}, a}$	Maximum voltage drop across the sense resistor during the acquisition a .
$V_{R \text{ min}, a}$	Minimum voltage drop across the sense resistor during the acquisition a .
x	Experiment index.
$\varepsilon_{\text{avg}, a}$	Relative error due to the use of the approximation $\tilde{P}_{\text{avg}, a}$ instead of $P_{\text{avg}, a}$.
$\varepsilon_{\text{char}, \mathbf{X}}$	Relative error on symbol \mathbf{X} (P_{avg} or t_{batt}) due to the character strings sent at regular intervals.
$\varepsilon_{\text{init}}$	Relative error on t_{batt} due to the initialization overhead ($\varepsilon_{\text{init}} = 0$).
ε_{res}	Relative error on t_{batt} due to the resolution interval t_{res} .
$\varepsilon_{\mathbf{X}}$	Relative error on symbol \mathbf{X} (e.g., $R, V, P_{\text{avg}}, t_{\text{batt}}$).

B Average power approximation

This appendix presents the development of the approximation $\tilde{P}_{\text{avg}, a}$ of the average power $P_{\text{avg}, a}$, which is used to establish Equation (10):

$$\begin{aligned}
\tilde{P}_{\text{avg}, a} &= \frac{\sum_{i=1}^N V_i}{N} \cdot \frac{\sum_{j=1}^N I_j}{N} = \frac{\sum_{i=1}^N \left(\sum_{j=1}^N V_i \cdot I_j \right)}{N^2} = \frac{\sum_{i=1}^N V_i \cdot I_i}{N^2} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j \neq i}^N V_i \cdot I_j \right)}{N^2} \\
&= \frac{\sum_{i=1}^N V_i \cdot I_i}{N^2} + (N-1) \cdot \frac{\sum_{i=1}^N V_i \cdot I_i}{N^2} - (N-1) \cdot \frac{\sum_{i=1}^N V_i \cdot I_i}{N^2} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j \neq i}^N V_i \cdot I_j \right)}{N^2} \\
&= N \cdot \frac{\sum_{i=1}^N V_i \cdot I_i}{N^2} + \frac{\sum_{i=1}^N \left(-(N-1) \cdot V_i \cdot I_i + \sum_{j=1, j \neq i}^N V_i \cdot I_j \right)}{N^2} \\
&= \frac{\sum_{i=1}^N V_i \cdot I_i}{N} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j \neq i}^N V_i \cdot (I_j - I_i) \right)}{N^2} \\
&= P_{\text{avg}, a} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j > i}^N V_i \cdot (I_j - I_i) \right) + \sum_{i=1}^N \left(\sum_{j=1, j < i}^N V_i \cdot (I_j - I_i) \right)}{N^2} \\
&= P_{\text{avg}, a} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j > i}^N V_i \cdot (I_j - I_i) \right) + \sum_{i=1}^N \left(\sum_{j=1, j > i}^N V_j \cdot (I_i - I_j) \right)}{N^2} \\
&= P_{\text{avg}, a} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j > i}^N V_i \cdot (I_j - I_i) \right) - \sum_{i=1}^N \left(\sum_{j=1, j > i}^N V_j \cdot (I_j - I_i) \right)}{N^2} \\
&= P_{\text{avg}, a} + \frac{\sum_{i=1}^N \left(\sum_{j=1, j > i}^N (V_i - V_j) \cdot (I_j - I_i) \right)}{N^2} \\
&= P_{\text{avg}, a} + \frac{\sum_{i=1}^{N-1} \left(\sum_{j=i+1}^N (V_i - V_j) \cdot (I_j - I_i) \right)}{N^2}
\end{aligned}$$

C Multimeter configuration

The following listings are the commands used to configure the HP 34401A multimeters [HP96a]. The initial number selects the multimeter (0 for V and 1 for V_R). The rest of the line is the actual command sent to the instrument, in *Standard Commands for Programmable Instruments (SCPI)* format. Commands are sent in listing order. The *initialization* section is sent once at the beginning of the experiment, while the *acquisition* section is sent for each acquisition.

Initialization

```

0 SYSTEM:REMOTE
1 SYSTEM:REMOTE
0 *RST
1 *RST
0 *CLS
1 *CLS
0 DISPLAY:TEXT "ITSY TEST: V"
1 DISPLAY:TEXT "ITSY TEST: I"
0 SYSTEM:BEEPER:STATE OFF
1 SYSTEM:BEEPER:STATE OFF
0 *IDN?
1 *IDN?
0 SYSTEM:VERSION?
0 FUNCTION "VOLTAGE:DC"
1 SYSTEM:VERSION?
1 FUNCTION "VOLTAGE:DC"
0 VOLTAGE:DC:RANGE:AUTO ON
1 VOLTAGE:DC:RANGE:AUTO ON
0 VOLTAGE:DC:NPLCYCLES 0.2
1 VOLTAGE:DC:NPLCYCLES 0.2
0 ZERO:AUTO ON
1 ZERO:AUTO ON
0 INPUT:IMPEDANCE:AUTO ON
1 INPUT:IMPEDANCE:AUTO ON
0 CALCULATE:FUNCTION AVERAGE
1 CALCULATE:FUNCTION AVERAGE
0 DATA:FEED RDG_STORE, ""
1 DATA:FEED RDG_STORE, ""
0 TRIGGER:SOURCE IMMEDIATE
1 TRIGGER:SOURCE IMMEDIATE
0 TRIGGER:DELAY MINIMUM
1 TRIGGER:DELAY MINIMUM
0 TRIGGER:COUNT 1
1 TRIGGER:COUNT 1
0 SAMPLE:COUNT 4150
1 SAMPLE:COUNT 4150

```

Acquisition

```

0 CALCULATE:STATE ON
1 CALCULATE:STATE ON
0 INITIATE ; CALCULATE:AVERAGE:AVERAGE?
1 INITIATE ; CALCULATE:AVERAGE:AVERAGE?
0 CALCULATE:AVERAGE:MINIMUM?
0 CALCULATE:AVERAGE:MAXIMUM?
0 CALCULATE:STATE OFF
1 CALCULATE:AVERAGE:MINIMUM?
1 CALCULATE:AVERAGE:MAXIMUM?
1 CALCULATE:STATE OFF

```

D Complete results

This appendix presents the average power P_{avg} , its lower and upper bounds $P_{\text{low bnd}}$ and $P_{\text{up bnd}}$, and the battery lifetime t_{batt} for all experiments. The product $P_{\text{avg}} \cdot t_{\text{batt}}$, representing the battery capacity, is also shown. This value is comparable to the manufacturer's rating of 640 mA · h. As expected, low-power experiments (i.e., sleep-mode and idle-mode benchmarks) make better use of the battery capacity than high-power experiments (i.e., run-mode benchmarks). The benchmarks were always executed in the following order:

1. Sleep.
2. Deep sleep.
3. Sleep, LCD.
4. MPEG-1, 206.4 MHz.
5. DECTalk, 73.7 MHz, low voltage (LV).
6. DECTalk, 73.7 MHz.
7. DECTalk, 206.4 MHz.
8. WAV, 59.0 MHz.
9. WAV, 206.4 MHz.
10. Idle, 59.0 MHz, low voltage (LV).
11. Idle, 59.0 MHz.
12. Idle, 132.7 MHz.
13. Idle, 206.4 MHz.
14. Sleep, daughter-card (DC).
15. Dictation, daughter-card (DC), 206.4 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	4.47	4.2 %	2.22	215	505.5	0.23 %	2.26
36	11	4.82	4.0 %	3.05	221	479.5	0.22 %	2.31
34	21	4.45	4.2 %	-0.78	206	514.5	0.23 %	2.29

Table 3: Benchmark results : deep sleep.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	7.46	3.0 %	5.20	217	302.5	0.24 %	2.26
36	11	7.56	2.9 %	5.78	221	304.5	0.24 %	2.30
34	21	7.19	3.0 %	5.92	206	318.5	0.24 %	2.29

Table 4: Benchmark results : sleep.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	10.6	2.4 %	8.46	239	211.5	0.29 %	2.23
36	11	10.6	2.4 %	8.88	245	214.5	0.29 %	2.28
34	21	10.5	2.4 %	8.69	228	218.5	0.29 %	2.29

Table 5: Benchmark results : sleep, daughter-card (DC).

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	26.3	1.6 %	23.9	237	85.5	0.61 %	2.25
36	11	26.3	1.6 %	24.4	239	86.5	0.60 %	2.28
34	21	25.9	1.6 %	24.4	214	88.5	0.59 %	2.29

Table 6: Benchmark results : sleep, LCD.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	57.0	1.9 %	52.4	395	38.8	0.45 %	2.21
36	11	54.7	1.9 %	50.5	377	41.1	0.45 %	2.25
34	21	54.5	1.9 %	50.7	373	41.8	0.45 %	2.28

Table 7: Benchmark results : idle, 59.0 MHz, low voltage (LV).

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	72.9	1.7 %	67.3	418	30.3	0.46 %	2.21
36	11	68.1	1.8 %	62.9	417	33.1	0.45 %	2.26
34	21	67.6	1.8 %	63.1	412	33.6	0.45 %	2.27

Table 8: Benchmark results: idle, 59.0 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	86.6	1.7 %	80.1	609	25.4	0.49 %	2.20
36	11	81.6	1.7 %	76.0	602	27.4	0.47 %	2.24
34	21	80.8	1.7 %	75.3	592	28.1	0.47 %	2.27

Table 9: Benchmark results: idle, 132.7 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	105	1.7 %	98	715	20.8	0.53 %	2.19
36	11	99	1.7 %	93	718	22.4	0.51 %	2.23
34	21	99	1.7 %	92	700	22.8	0.51 %	2.25

Table 10: Benchmark results: idle, 206.4 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	283	1.8 %	149	671	7.47	0.58 %	2.11
36	11	276	1.8 %	146	664	7.82	0.57 %	2.16
34	21	276	1.8 %	144	639	7.96	0.56 %	2.20

Table 11: Benchmark results: WAV, 59.0 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon_{P_{\text{avg}}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon_{t_{\text{batt}}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	316	2.3 %	180	990	6.62	0.63 %	2.09
36	11	307	2.4 %	174	980	6.97	0.61 %	2.14
34	21	308	2.3 %	174	964	7.04	0.60 %	2.17

Table 12: Benchmark results: WAV, 206.4 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon P_{\text{avg}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon t_{\text{batt}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	355	1.7 %	53.3	638	5.97	0.57 %	2.12
36	11	352	1.7 %	51.1	598	6.13	0.56 %	2.16
34	21	350	1.7 %	51.8	571	6.24	0.56 %	2.18

Table 13: Benchmark results: DECTalk, 73.7 MHz, low voltage (LV).

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon P_{\text{avg}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon t_{\text{batt}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	402	1.6 %	70.1	634	5.20	0.63 %	2.09
36	11	397	1.6 %	65.3	645	5.37	0.61 %	2.13
34	21	393	1.6 %	65.4	627	5.47	0.60 %	2.15

Table 14: Benchmark results: DECTalk, 73.7 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon P_{\text{avg}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon t_{\text{batt}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	405	2.9 %	97	967	5.13	0.63 %	2.08
36	11	401	2.9 %	91	944	5.29	0.62 %	2.12
34	21	397	2.8 %	91	940	5.45	0.60 %	2.16

Table 15: Benchmark results: DECTalk, 206.4 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon P_{\text{avg}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon t_{\text{batt}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	762	2.8 %	107	1250	2.61	0.65 %	1.99
36	11	760	2.8 %	100	1260	2.66	0.64 %	2.02
34	21	750	2.8 %	99	1230	2.75	0.62 %	2.06

Table 16: Benchmark results: dictation, daughter-card (DC), 206.4 MHz.

Itsy no.	Batt. no.	P_{avg} [mW]	$\varepsilon P_{\text{avg}}$	$P_{\text{low bnd}}$ [mW]	$P_{\text{up bnd}}$ [mW]	t_{batt} [h]	$\varepsilon t_{\text{batt}}$	$P_{\text{avg}} \cdot t_{\text{batt}}$ [W · h]
35	17	830	3.6 %	229	1400	2.35	0.64 %	1.95
36	11	821	3.7 %	223	1380	2.42	0.63 %	1.99
34	21	813	3.5 %	222	1380	2.50	0.61 %	2.03

Table 17: Benchmark results: MPEG-1, 206.4 MHz.