

Energy Efficient Sensor Node Implementations

Jan R. Frigo, Eric Y.
Raby, Sean Brennan
Los Alamos National
Laboratory
ISR-3, Los Alamos, NM
MSD440, Los Alamos, NM
{jfrigo, raby,
brennan}@lanl.gov

Edward Rosten
Dept. of Engineering
University of Cambridge, UK
er258@cam.ac.uk

Christophe Wolinski,
Charles Wagner,
Francois Charot
IFSIC/IRISA
University of Rennes, France
{wolinski, wagner,
charot}@irisa.fr

Vinod K. Kulathumani
Dept. of Computer Science
and Electrical Engineering
West Virginia University
Vinod.Kulathumani@mail.wvu.edu

ABSTRACT

In this paper, we discuss a low power embedded sensor node architecture we are developing for distributed sensor network systems deployed in a natural environment. In particular, we examine the sensor node for energy efficient *processing-at-the-sensor*. We analyze the following modes of operation; event detection, data acquisition, and data processing using low power, high performance embedded technology such as specialized embedded DSP processors and low power FPGAs at the sensing node. We use compute intensive sensor node applications: an acoustic vehicle classifier (frequency domain analysis) and a video license plate identification application (learning algorithm). We report performance and energy for these applications and discuss the system architecture design trade offs.

Categories and Subject Descriptors

C.3 [Computer Systems Organization]: Special-Purpose and Application-Based Systems, Real-time and embedded systems

General Terms

Algorithms, Design, Performance

Keywords

FPGA, DSP, Distributed Sensor Network (DSN), seismic, acoustic, video, vehicle classification

1. SENSOR NODE ARCHITECTURE

In this section, we give an overview of our sensor node architecture. Section 2 describes the vehicle classifier and the

Copyright 2010 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.
FPGA 2010 February 21–23, 2010, Monterey, CA USA
Copyright 2010 ACM 978-1-60558-911-4/10/02 ...\$10.00.

license plate identification algorithms used for benchmarking. In Section 3 the energy and performance results are given for various FPGA families and embedded processors.

Our node architecture [4] consists of a carrier board with the following: an ARM mezzanine board¹, an embedded GPS module, a wireless chip, and four sensor interface connections. The ARM mezzanine board is designed around the NXP LPC3180 chip and can be used for rapid prototyping or as a high performance co-processor. It can be removed if a co-processor is not needed, i.e. if the node is a relay or it can be replaced with an application specific processor board. In our system, sensor boards connect directly to the carrier board and performs the following functions; data acquisition, data processing and power management. The sensor board design consists of an embedded processor or FPGA, an A/D converter and signal conditioning circuitry.

2. APPLICATIONS

2.1 Vehicle Classifier

The vehicle classification system was developed using seismic and acoustic sensor data to classify vehicles as they approach a specific region of the roadway [3]. The vehicle classification and detection algorithms are described in [4]. For vehicle detection a level-two Haar Wavelet is used per Eqn. 1 where $h(i)$ denotes the 32 coefficients of the Haar wavelet transform corresponding to a particular frequency band. The energy estimate and the variance of the energy estimate is computed using Eqn. 2 and Eqn. 3. A variance threshold is used for vehicle detection. Once a vehicle is detected, real-time acoustic data is processed using a 16-bit, 512-point, integer FFT until a classification is determined.

$$amp(t) = (1/32) * \sum_{i=0}^{31} h(i) \quad (1)$$

$$mean(t) = (1/20) * \sum_{i=0}^{19} amp(t - i) \quad (2)$$

¹Phytec phyCOREARM9/LPC3180

$$var(t) = (1/20) * \sum_{i=0}^{19} amp(t-i) - mean(t)^2 \quad (3)$$

Figure 1 shows the pipelined FPGA architecture corresponding to the seismic detection algorithm. It is composed of four processing modules, $PM0 - PM3$ and one control module, $CM0$. The $PM0$ module saves and scales every 10th sample from the input data stream. The $PM1$ module executes the level-two wavelet transformations. Modules $PM2$ and $PM3$ compute the mean and variance of the wavelet coefficients according to Eqns. 2 and 3. For the acoustic processing algorithm, we examine 512-pt integer FFT cores generated from the Xilinx ISE, and ACTEL Libero IDE FPGA development tools. Table 1 and Table 2 give the performance and energy results for the seismic detection algorithm (Eqns 1 - 3) and the 512-pt integer FFT on reconfigurable and embedded architectures.

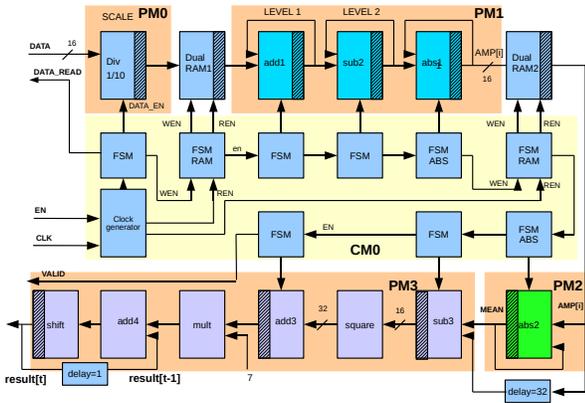


Figure 1: Level-2 Haar wavelet, mean and variance FPGA implementation

2.2 License Plate Identification

The license plate identification system extracts license plate information from a moving vehicle on a roadway using video and magnetometer data [4], [3]. The three processing modes for the license plate identification application are: magnetometer detection, image capture and license plate pixel identification.

First, the magnetometer detection algorithm signals an event by continuously acquiring data (at 1 kHz) and computing the error as the sum of squared differences between the magnetometer input and the filter state. The signal is filtered with an IIR filter per Eqn. 4 where A is a constant equal to 2^{-8} . The 3-axis magnetometer provides the values x , y and z (with the filtered measurements being xi , yi and zi respectively). The error is calculated per Eqn. 5. When the error passes a set threshold, image capture is initiated and a selection routine determines the *best* image to process for license plate identification.

$$filter_{t+1} = filter_t - A * filter_t + A * input_t \quad (4)$$

$$error = (x - xi)^2 + (y - yi)^2 + (z - zi)^2 \quad (5)$$

The license plate identification algorithm as described in [4] works by applying a classifier to every pixel in an image to create a rough segmentation of the license plate, if it exists. From this, the bounding box of the license plate is found, and that section of the image is then resampled to a fixed size. The resampled image is then PNG compressed and sent over the network to a base station computer. Table 3 gives benchmarking results for the license plate identification algorithm using two embedded processors and the Crossbow Stargate (Intel XScale).

3. RESULTS

In this section we examine total energy utilized for each application. We analyze computational processing, event detection and data acquisition modes.

The FPGA energy results herein are derived from the Altera (Quartus II), Xilinx (Xpower), and ACTEL (SmartPower) development tools and our field experiment data per the frequency of the routed hardware design. The results² in Table 1 and Table 2 show energy consumed, total power, execution time, throughput, and run-time frequency for each application. The total power is a combination of both quiescent and dynamic power. The embedded processor energy results herein are actual measurements taken from benchmarking.

This study includes the following reconfigurable and embedded architectures: Xilinx, Altera and ACTEL FPGAs³, a DSP processor⁴, two embedded processors⁵, and Crossbow's Mica2⁶ and Stargate⁷ processors. We include the Crossbow devices in our benchmark results as both are frequently used for sensor network implementations by the research community.

3.1 Processing

Vehicle Classifier: The results from Table 1 and Table 2 show an expected trend, the specialized, low power, embedded architectures show more energy efficiency. In this case, the FPGAs have the lowest energy utilization for compute-intensive data processing. If we consider the wavelet transformation, mean and variance calculations, there is approximately a 424x energy savings using the Igloo FPGA over the LPC3180, the most efficient embedded processor. For the 512-pt integer FFT in Table 2, the FPGA devices are 12 (Igloo) to 38 (Sparten3) times more energy efficient than the most efficient embedded processors (Blackfin and optimized DPS). Comparing one FFT computation on the Blackfin processor (19 μ J) with the Proasic3 (1.77 μ J) device, we can process 2 times the amount of data with a 10 times improvement in energy utilization.

²

$$throughput(MSPS) = n(samples)/executiontime(s)$$

$$energy(J) = measurepower(J/s) * executiontime(s)$$

$$executiontime(s) = n(cycles)/clockfrequency(cycles/s)$$

³Virtex4 XC4VLX15, Spartan3 XC3S400, Stratix II EP2S60, CycloneII EP2C35F, Igloo AGL1000V5)

⁴Texas Instrument's TMS320C5510

⁵NXP LPC3180 and Analog Devices Blackfin ADI BF537 processor.

⁶ATMEL ATmega128

⁷Intel XScale

	time	pwr	freq	thru	energy
	μs	mW	MHz	MSPS	μJ
Igloo	5.5	5.95	23.16	23	0.033
CycloneII	1.3	164	100	98	0.21
Stratix II	1.3	731	100	98	0.95
Virtex4	1.22	288	105	105	0.35
Spartan3	2.28	120	56	56	0.27
Mica2	1077	60	4	0.119	65
DSP	145	262	200	0.882	38
LPC3180	42.8	330	208	3.0	14
Blackfin	21	1056	500	6.1	22

Table 1: Seismic Processing Energy Utilization (Level 2 Haar wavelet, mean, variance)

	time	pwr	freq	thru	energy
	μs	mW	MHz	MSPS	μJ
Igloo	8.4	185	61	61	1.55
Proasic3	6.53	272	78.3	78.3	1.77
Virtex4	5.12	304	100	100	1.55
Spartan3	5.12	97	100	100	0.5
Stargate	867	1850	400	0.59	1604
DSP	3000	227	200	0.17	681
DSP opt	60	339	200	8.5	20
LPC3180	877	400	208	0.58	351
Blackfin	653	1175	500	0.78	767
Blackfin ¹⁰	139	1175	500	3.7	163
Blackfin ¹¹	18.8	1007	500	27	19

Table 2: Acoustic Processing Energy Utilization (512-pt int FFT)

License Plate Identification: The license plate identification processing algorithm performs the following functions: preprocessing, detection of license plate pixels, centroid, resampling, and compression. The energy estimates in Table 3 show the total energy for these functions. In this application, the algorithm is a decision tree. The tree is applied independently at every pixel of the image. The decision tree was used in this implementation for computational speed on the CPU, thus, the benchmark results are for the Stargate (XScale) and two embedded processors, the LPC3180 and the Blackfin. The most energy efficient processor is the LPC3180 at 0.185 J (Table 3).

The other processing modes for the license plate identification application are magnetometer detection and image capture. The magnetometer detection mode will be discussed in Section 3.2. Image capture is insignificant in terms of energy utilization compared to the pixel identification algorithm due to anticipated fast processing time, i.e. assuming one or two frames are captured in approximately 100 ms. Using maximum active MPU mode on the LPC3180 at 80 mA, 1.2 V [6], the estimated energy is 9.6 mJ.

3.2 Data Acquisition and Event Detection

Seismic Detection: For the vehicle classifier system, the seismic algorithm signals an event by continuously acquiring and thresholding seismic data at 100 Hz. When a vehicle approaches the seismic detection algorithm is computed (wavelet, mean and variance) and acoustic classifica-

	time	pwr	freq	thru	energy
	s	W	MHz	KSPS	J
Stargate	0.675	1.575	400	455	1.06
LPC3180	0.608	0.305	208	503	0.185
Blackfin	0.534	1.148	500	574	0.61

Table 3: License plate identification energy utilization

tion processing is initiated. The system is sensing continuously during *active* time periods. During *rest* time periods, the system is *sleeping* and *waking up* to acquire data and check for events. The question arises, if an FPGA is most energy efficient for computation, is it comparable to a microcontroller for wake-up timing, sleep mode power (considering leakage current) to control event detection and data acquisition? In addition, should power management be performed by the FPGA for energy efficiency? In the following, we will analyze the Igloo FPGA, Mica2 and the LPC3180 with respect to these operating modes—wake-up timing, sleep, idle and active (event detection) modes per the seismic detection algorithms.

The Igloo FPGA family is optimized for ultra-low power, embedded applications like seismic event detection. These devices have Flash*Freeze technology that enables fast switching from ultra-low power modes. Power management does not require extra components to turn off I/Os or clocks and retains design, SRAM content and registers. Wake up timing is reported as 1 μs [1]. For this application, the quiescent Flash*Freeze mode leakage power is 114 μW (Table 4) with all voltages (including the core voltage, $V_{cc} = 1.5 V$) on and all clocks and I/Os off. For sleep mode with only the core voltage on and all other voltages, clocks and I/Os off, power usage is 10.8 μW . Idle power is 134 μW . Run time operating frequency for this application is 23.16 MHz.

The Mica2 is commonly used in the sensor network community for deployed DSN applications [5]. It has an AT-MEL ATmega128 8-bit RISC microcontroller with six sleep modes and supports many standard peripherals. We compute stand-by power to be 202.5 μW (75 μa , 2.7 v with a 4MHz XTAL) [2]. Wake up for stand-by mode is reported as 1.5 μs . The highest operating speed is 16 MHz.

The NXP LPC3180 device is a more powerful processor that also could be used for seismic vehicle detection, especially if classification utilized seismic data. During continuous floating point operation power consumption of the mezzanine board is measured to be approximately 330 mW. The low power modes are as follows: direct RUN is 7 mA at 13 MHz (slow clock), and STOP mode⁸ is 450 μW (500 μa at 0.9 v). For wake-up timing, typical values for ARM CPUs are less than 0.5 ms. Leakage current for the LPC3180 is reported as 3 μa [6]. The highest operating speed is 208 MHz.

If we compare the average power results shown in Table 4, where *on time* is determined as *wake up time + compute time* (from Table 1), it is clear that the ultra-low power FPGA has an advantage over the other devices since the performance is 23.13 MHz compared to a 100 Hz sampling rate, we can process at approximately 10k times the sam-

⁸This power is achievable only if the chip is isolated, without external memory or a Linux OS running, etc.

	wake-up time	sleep mode	on time	ave pwr
	s	W	s	W
Igloo	1 μ	114 μ	5.3 μ	114 μ
Cyclone II	1 μ	80 $m\ddagger$	1.1 μ	80 m
Stratix II	1 μ	624 $m\ddagger$	1.1 μ	624 m
Mica2 (4 MHz)	1.5 $\mu\§$	202.5 μ	1.08 m	267 μ
LPC3180	0.5 m	450 μ	542.8 μ	629 μ

Table 4: Average power for seismic event detection (§wake up from standby, ‡static power)

pling rate. Since the device has a 114 μ W Flash*Freeze mode (low power sleep mode) in which it is operating for over 99.99% of the time, the power profile of this device has a 2.3 to 5.5 times average power saving over the embedded microcontrollers (see Table 4).

In addition, we notice that other FPGA devices have excellent performance such as Cyclone II and Spartan3, but these devices do not have ultra-low power sleep modes—a necessary requirement for deployed sensor network applications. Since FPGAs have 100’s of I/O, and can process in parallel, the FPGA can compute vehicle detection and classification, acquire sensor data and manage power. Thus, there is no need for a secondary microcontroller at the sensor board which reduces the sensor node footprint as well as improves the overall power budget for the sensor board.

Magnetometer Detection: For the license plate identification system, the magnetometer algorithm signals an event by continuously acquiring data at 1k Hz and computing the error as the sum of squared differences between the magnetometer input and the filter state. At this sampling rate, use of the LPC3180 is not practical since the wake-up time is too slow (0.5 ms). Thus, the magnetometer detection algorithm is well suited for implementation on an FPGA and we expect the energy utilization to be insignificant compared to license plate identification processing.

	active pwr	on time	energy
	mW	μs	μJ
seismic acq	0.134	134	0.018
seismic detect	5.95	5.5	0.033
acoustic acq	0.216	616	0.132
512-pt fft	185	8.4	1.55

Table 5: Vehicle Classifier Event Energy (Igloo)

	active pwr	on time	energy
	mW	s	mJ
image capture	96	0.10	9.6
lp ident alg	305	0.61	185

Table 6: License Plate Id Event Energy (LPC3180)

4. CONCLUSION

In this paper we investigate event detection, data acquisition, data processing modes at the sensing node using two compute-intensive applications. Our node architecture aims

to keep the data transfer, both intra-module and to the network, at a minimum by processing data in real-time on the sensor board. In our system, we have flexibility to choose the most suitable embedded technology for the task. The master controller for event detection, data acquisition and power management is the FPGA or microcontroller on the sensor board. Memory required to store events is minimized since raw sensor data is significantly reduced. Sensor data is reduced by *processing-at-the-sensor* from 80 KB to 1 byte for the vehicle classification application and from 300 KB to approximately 3 KB for the license plate identification application. Finally, a significant savings in energy is realized over COTS implementations of these applications. The vehicle classifier application utilizes a total of 1.73 μ J on the Igloo device (Table 5) and the license plate identification application utilizes a total energy of 0.195 J on the LPC3180 processor (Table 6). Using COTS hardware [4], the vehicle classification application took 10 seconds to process using a total of 18 to 22 J and the license plate identification application took 5.61 seconds and used 14.5 J.

There are several noteworthy insights as a result of this study. Sleep mode dominates the total energy utilization profile for both of these applications, thus, the processor on the sensor board must have fast wake-up timing and ultra-low power sleep modes. Operations such as data acquisition and power management can be handled by the FPGA to save power and reduce the node footprint. In terms of data processing on the sensor node, the FPGAs have excellent performance and energy savings compared to the embedded processors benchmarked herein. The FPGA results in Tables 1 and 2 show at least a 12 (Igloo) to 38 (Spartan3) times energy savings over the most efficient embedded processor (Blackfin) for the FFT algorithm and a 424 times savings for the seismic detection algorithms using the Igloo device compared to the LPC3180 processor.

5. ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy/NNSA and Los Alamos National Laboratory funds under LANS, LLC Contract No. W-7405-ENG-36. This document is approved for public release under LAUR-09-05907.

6. REFERENCES

- [1] ACTEL. Igloo. <http://www.actel.com/products/igloo>, 2009.
- [2] ATMEL. Atmega128l. <http://www.atmel.com/dyn/products>, 2009.
- [3] J. Frigo, V. Kulathumani, S. Brennan, E. Rosten, E. Esch, D. Jackson, P. Majerus, A. Warniment, A. Mielke, and M. Cia. Radiation detection and situation management by distributed sensor networks. In *SPIE Proceedings on Defense, Security and Sensing*, 2009.
- [4] J. Frigo, V. Kulathumani, S. Brennan, E. Rosten, and E. Raby. Sensor network based vehicle classification and license plate identification system. In *IEEE INSS*, 2009.
- [5] L. Gu and et. al. Lightweight detection and classification for wireless sensor networks in realistic environments. *SenSys 2005*, Nov. 2005.
- [6] N. Phillips. Lpc3180 processor. <http://www.standardics.nxp.com/products/lpc3000/lpc3180>, 2009.