

A Low-Power Wireless Camera System

Anantha Chandrakasan, Abram Dancy, James Goodman, Thomas Simon
Department of EECS,
Massachusetts Institute of Technology, Cambridge

ABSTRACT

This paper describes the system design of a low-power wireless camera. A system level approach is used to reduce energy dissipation and maximize battery lifetime. System properties such as the network configuration and data statistics are exploited to minimize computational switching. Embedded power supplies systems are also used to minimize energy dissipation under varying temperature, process parameters and computational workload. Since the camera operates in a burst mode with long idle periods, emphasis must be placed on reducing system standby power dissipation.

1. Introduction

This paper describes the design considerations for an ultra low-power wireless camera (Figure 1). The camera transmits compressed video data over a wireless link to a fixed base station. The data rate is variable up to a maximum of 1Mbps (the image sensor has a spatial resolution of 256x256 quantized to 8 bits/pixel). Many of the design issues faced in the context of our wireless camera are common to those in other wireless applications. Total system energy (computation and communication) averaged over the normal operating conditions of the device should be minimized to maximize battery lifetime. The system should also be designed to service time varying data rates and quality of service requirements; embedded power supplies, which adapt supply voltages on demand, can save significant power in such systems.

The main application specific issues of our camera have to do with the asymmetry between a camera and the receiving base-station, which may be communicating with

several cameras. The base-station is assumed to not be battery operated and this creates opportunities to design a system with skewed computational burdens, favoring the low-power cameras.

2. Video Compression

The bandwidth over the wireless link is limited and therefore data must be compressed before transmission. The image compression algorithm/architecture must also be capable of optimizing system wide power for widely differing output bit rates. During, possibly long, periods of no motion in the image source, the system power is determined by the computational cost of determining there is no movement, the operation of the sensor and converter, and the standby losses of all modules and power supply. Besides minimizing standby losses in individual modules through architectural and circuit techniques, system power may be reduced by the image compression module by adaptive control of frame rate, bit resolution, and feedback regarding optimal operating voltages to the power supply.

The compression algorithm is based on Shapiro's image compression work using zero-tree coded subband decompositions [1]. A multi-resolution wavelet decomposition of a frame is computed. The coding scheme for the quantized coefficients takes the expected distribution across subbands into account in order to achieve good compression. This image compression algorithm was chosen because of its excellent compression and low computational cost. In addition, the algorithm produces an embedded output stream. This means that more highly compressed (and poorer quality) codings of an image are the prefix for refined codings. In other words, the bits that represent an image are sent in order of importance, according to some metric. Hence, there is a convenient knob to turn in order to trade-off compression for image quality on the fly.

Power dissipation can be reduced at the architecture and circuit levels as well as at the algorithmic level. Typically, this involves parallelizing the execution of an algorithm to maintain throughput, while slowing the operation of circuits to reduce total power. A highly parallel SIMD architecture is used to operate the processor at a supply voltage of 1V. Factors that determine the power dissipation of the resulting implementation include the number and complexity of logical/arithmetic operations, size and

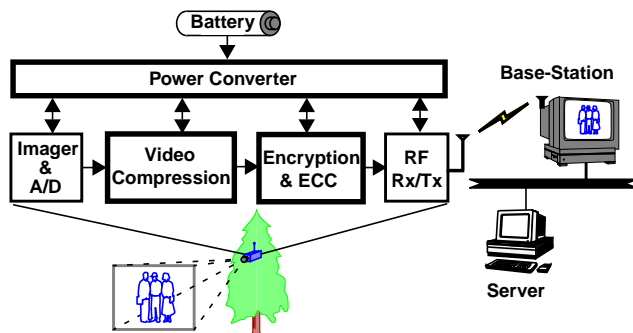


Figure 1. MIT Wireless Video Camera.



Figure 2. Image at various compression ratios.

usage pattern of memory, and amount of non-local communication. Each processing element has embedded memory in order to minimize the communication costs and the energy dissipated in memory access.

An adaptive arithmetic coder is used in the system. Arithmetic coding results in 15 to 20 percent higher compression compared to Huffman coding of similar complexity, so the use of arithmetic coding is desirable. However, computing an adaptive arithmetic code involves divisions by non-constant values, which is an expensive operation. We found that using an approximating coder, which rounds probability estimates so that all divisions are by powers of 2, results in < 1 percent loss of compression.

The choice of filters is an important consideration for low power systems. Typically, shorter filters give poorer compression performance but involve fewer arithmetic operations as well as simpler ones. The loss in compression is often small, favoring the choice of short filters. Due to the asymmetric computational constraints on coder and decoder, asymmetric filters are a good choice in this application.

During periods of little motion, simple frame differencing performs quite well compared to fully motion compensated differencing. Even with large amounts of motion, simple differencing still outperforms intra-frame coding of single frames. Performing motion estimation at the encoder would dramatically increase the computational complexity and power. However, real motion vectors are typically predictable from previous frames' vectors, over many frames. This presents the possibility of coding a few frames without motion compensation, having the base-station compute motion vector guesses from the first frames, feeding back these guesses to the camera (without too much latency) and performing only correction steps to these vectors at the encoder for succeeding

frames [2]. Figure 2 shows the image coding results for various compression ratios.

3. Data Encryption

In many applications, such as the wireless camera, it is desirable to design digital processors that allow a trade-off between the quality of service (QoS) provided and the energy consumed to process a sample. This allows the user to evaluate the application's requirements and set the desired quality while minimizing the energy consumption. We have developed an energy scalable encryption processor where the level of security (i.e., quality) and energy consumed to encrypt a bit can be traded-off dynamically based on demand. Since transmitted data streams can often be partitioned into different priority levels, an energy scalable processor ensures that important information is adequately protected, while sacrificing some security for low priority data in order to reduce the total system energy.

The energy scalable encryption processor in this work is based on a variable-width quadratic residue generator (QRG). The QRG is a cryptographically-secure pseudo-random bit generator that is based upon the work in [3]. The QRG operates by performing repeated modular squarings. The modular squaring is performed using an algorithm based on Takagi's iterated radix-4 algorithm [4] which requires $(\log_2 Q)/2$ iterations to compute the result $P = X \cdot Y \text{ mod } Q$. The least significant $\log_2 \log_2 Q$ bits of each result can be extracted and used as a strong reproducible pseudo-random source for applications such as a stream cipher or key generator.

Energy scalable computing requires dynamically reconfigurable architectures that allow the energy consumption per input sample to be varied with respect to quality. In the case of the QRG, the quality scales sub-exponentially with the modulus length, while the energy

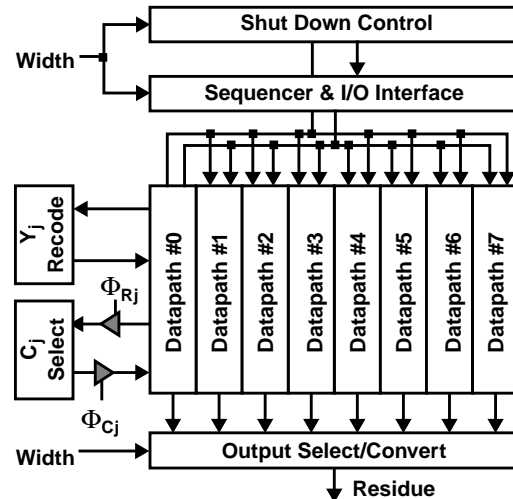


Figure 3. Energy scalable encryption processor [5].

consumption scales polynomially. A fully scalable QRG architecture was developed where the width ($w = \log_2 Q$) can be reconfigured on the fly to range from 64 to 512 bits in 64 bit increments (Figure 3) [5]. The design makes extensive use of clock gating to disable unused portions of the QRG. Hence the switched capacitance of the QRG is minimized and energy scalability is achieved.

Further energy/security scalability can be achieved through the use of an adaptive supply [6]. Rather than designing a system with a static supply to meet a specific timing constraint under worst case conditions, it is more energy efficient to allow the voltage to vary such that the timing constraints are just met at any given temperature and operating conditions. In the encryption processor, when operating at a reduced width, the number of cycles required per multiplication is reduced and therefore the supply voltage can be reduced for a given throughput. The supply is varied using an embedded custom DC/DC converter. The use of an adaptive supply enables us to substantially reduce the energy consumption as the multiplier width is varied (Figure 4). Figure 5 shows a plot of security (measured in MIPS-years) as a function of energy per bit. This plot was obtained by varying the bit-

width and supply. Figure 6 shows a die photo of the scalable encryption processor with embedded power supply.




Figure 6. Die photo of encryption processor.

4. High Efficiency DC-DC Conversion

In portable systems, such as the wireless camera, the electronic circuits can be designed to operate over the range of the voltages supplied by the battery over its discharge cycle. However, adding some form of power regulation can significantly increase battery life, since it allows circuitry to operate at the “optimal” supply voltage from a power perspective. As seen from the previous section, the optimal power supply voltage can change dynamically based on the throughput or quality requirements. Therefore, the DC-DC converter must be designed to handle widely varying power supply voltages and power levels.

Figure 7 shows a block diagram of the DC-DC converter. The converter operates by creating a pulse width modulated signal of some duty cycle at the input to the LC filter, whose average value is the desired output voltage. External passive filtering is used to filter the PWM signal, creating a DC voltage with some tolerable value of ripple.

In order to provide reasonable efficiencies for the low supply voltages present in low power digital systems, power converters must incorporate synchronous rectification (i.e., active power devices are used to replace diodes). A drawback of synchronous rectification is that without explicit monitoring of the output current and control of the synchronous rectifier, the circuit will not enter

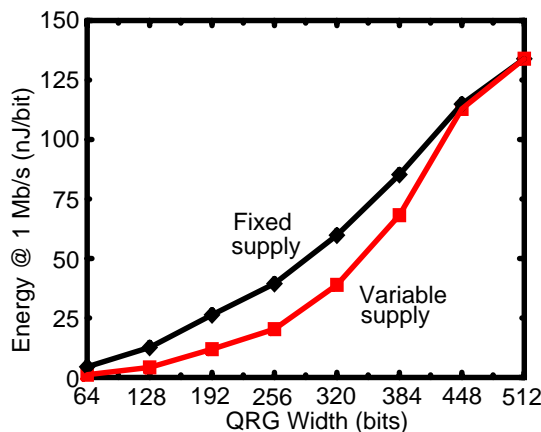


Figure 4. Power dissipation for fixed and variable supply.




Figure 5. Security vs. Energy.




Figure 7. Architecture of the DC-DC Converter.

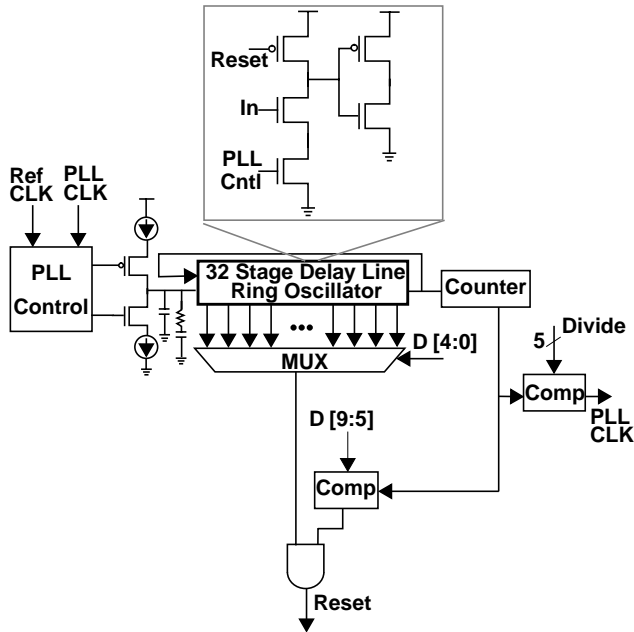


Figure 8. PLL based PWM generation block.

discontinuous mode at light loads. The resulting ripple current in the inductor will cause resistive losses that will reduce efficiency at light loads. Hence, the ability to create a “turn-off” signal for the synchronous rectifier could be an important feature for a low power controller.

The compensation network for the output of the power converter is a variable gain integral controller. A reference value (in a digital form) is subtracted from the A/D or delay measurement, and the difference is scaled in an array multiplier stage. The product is then subtracted from the previous duty cycle command to produce the next duty cycle command. The internal representation of the duty cycle is 12 bits, and the 10 MSBs are passed to the PWM stage to create the output. The compensation sample rate is programmable; the sample rate is primarily limited by the A/D conversion time. The reference value and gain for the output and other configuration registers are fully programmable through a bidirectional two wire serial interface.

A PLL based approach (Figure 8) is used for generating the PWM signal. A 32 stage delay line forms the basis for the pulse width modulation stage. The delay line is configured as a ring oscillator, which is phase locked to a reference clock. A divider allows the ring oscillator frequency to be set between 2 and 32 times faster than the reference frequency. The taps of the delay line then divide the input clock period into between 64 and 1024 equal increments. The taps of the delay line are sensed by two 32 to 1 multiplexers, one for each of the output PWM signals. The rising edge of the reference clock sets the PWM signals high. A PWM signal is set low when a pulse arrives at the tap of the delay line selected by its multi-




Figure 9. Low voltage modified Widlar 100nA current source, using MOS devices in subthreshold. This circuit generates bias voltages for the PLL charge pump. All dimensions in μm .




Figure 10. Transient response of filtered output voltage to changing digital reference commands ($f_{SW}=1\text{MHz}$).

plexer for the Nth time, where N represents the 5 MSBs of the 10 bit duty cycle command.

The delay of the delay line is controlled by adjusting the gate signals on starvation-type NMOS devices. The gate control signal controls the speed of the positive going edge at the output of each buffer. The control node is charged up and down using a current source. The biasing for the current source is generated on chip with a MOS Widlar current source (Figure 9). The compensation network for the PLL control node is also implemented on chip with poly-poly capacitors and a poly-2 resistor.

The efficiency of the converter was measured to be > 90% while delivering a load of 1mA at 1V. Figure 10

shows the transient response of the DC-DC converter in response to a change in desired operating voltage.

5. Conclusion

Low-power design requires a system level methodology that explicitly considers computation and communication costs. Extremely low power operation can be achieved in digital circuits by aggressively scaling of the power supply voltage. In many cases, the supply has to be adaptive to meet time varying QoS or data rate requirements. In order to change the voltage dynamically, efficient DC-DC conversion circuitry must be designed for widely varying voltages and loads.

Acknowledgments

This work is sponsored by DARPA under contract #DAAL-01-95-K3526.

References

- [1] J. Shapiro, Embedded image coding using zerotrees of wavelet coefficients, IEEE Transactions on Signal Processing, Dec. 1993.
- [2] W. Rabiner and A. Chandrakasan, "Network-Driven Motion Estimation for Wireless Video Terminals," IEEE Transactions on Circuits and Systems for Video Technology, pp. 644-653, Aug. 1997.
- [3] L. Blum, M. Blum, and M. Shub, "A Simple Unpredictable Pseudo-Random Number Generator," SIAM Journal on Computing, v. 15, n. 2, pp. 364-383, 1986.
- [4] N. Takagi, "A Radix-4 Modular Multiplication Hardware Algorithm for Modular Exponentiation", IEEE Transactions on Computers, pp. 949 - 956, Aug. 1992.
- [5] J. Goodman, A. P. Chandrakasan, "A 1Mb/s Energy/Security Scalable Encryption Processing Using Adaptive Width and Supply," IEEE International Solid-state Circuits Conference, pp. 110-111, Feb. 1998.
- [6] V. Gutnik, A. Chandrakasan, "Embedded Power Supply for Low-Power DSP," IEEE Transactions on Very Large Scale Integration (VLSI) Systems, pp. 425-435, Dec. 1997.