

Adaptive Low-Power Analog/Digital Converters for Wireless Sensor Networks

Dietmar Schroeder

Department of Microelectronics
University of Technology, Hamburg, Germany
d-schroeder@tu-harburg.de

***Abstract** — The paper addresses the problem of power consumption of sensor nodes in a wireless network. An integrated low-power analog/digital converter (ADC) is presented that is particularly suited for wireless sensor applications. The converter makes use of information theoretic redundancy in the input signal for reducing the conversion workload and performing data compression on-the-fly during conversion (entropy-coding A/D converter). Thus, energy is saved both in signal conversion and transmission. Experimental results from a prototype chip are presented.*

The converter is especially suitable for sensor networks that maintain a global data model. This is further illustrated on an exemplary scenario of distributed wave detection. It is shown that sensor signal detection and acquisition in this type of applications can be carried out very efficiently with entropy-coding converters used in the sensor nodes.

1 Introduction

Since the sensor nodes in future wireless sensor networks are expected to work autonomously with little energy supply for a long time, the problem of low power consumption of the nodes is an important issue. Thus it is advisable to include possibilities on all design levels to reduce the power consumption in the system design process. While the power consumption of the digital parts of the sensor nodes benefit from the advances of semiconductor technology, the same statement cannot be made for the analog and radio parts [1]. This paper addresses in particular the reduction of power consumption of the analog-to-digital conversion by pointing out the particular suitability for wireless sensor networks of our recent A/D converter architecture.

We have developed a novel circuit architecture for integrated analog/digital converters (ADCs) that allows to exploit the information-theoretic redundancy of the input signal for increasing the efficiency of operation and reducing the power consumption of the converter [2]. During analog/digital conversion the ADC performs data compression at the same time, without extra cost. This fact reduces the necessary bandwidth for radio transmission of the sensor values, which in turn diminishes the workload of the digital post-processing and the radio transmitter as well. Details on

the basic concepts of the ADC and experimental results from a prototype chip implementation [2] are reported in Section 2.

Section 3 is intended to illustrate the benefit that can be obtained in a wireless sensor network from this type of ADC. For this purpose, we consider a scenario where a network of sensors is distributed over some area of interest. The sensors measure a certain physical quantity, for instance pressure. We expect the occasional occurrence of waves of this quantity. The goal for the network is to monitor the area and to give alarm in case of the occurrence of a wave, and in this case to observe the wave with high precision. It is discussed that employing a global model of the sensor field, a considerable amount of energy can be saved if entropy-coding ADCs are utilized in the sensor nodes. Conclusions are drawn in the final Section 4.

2 Entropy-coding Analog/Digital Converter

In wireless sensor networks it is mandatory that the analog/digital converters in the sensor nodes have a power consumption as small as possible. In addition, it is beneficial for the battery lifetime of the node if we compress the sensor data prior to transmission over the network. Usually, data compression is done using one or another form of entropy-coding. However, if it is possible to compress the sensor data after digitization at all, this means that the ADC has created unnecessarily many bits and consumed unnecessarily much energy. Thus, an important increase in efficiency can be obtained if we combine the A/D-conversion and the entropy-coding into a single process, because the operational effort of the ADC then can be adapted to the true information rate of the sensor signal. This idea has been carried out into the development of an integrated entropy-coding analog/digital converter (ECOC) [2].

According to information theory, the information flow is the smaller the larger the pre-knowledge of the receiver about the information source is. The measure for the uncertainty of the incoming signal is the entropy [3]. If the receiver knows much about the incoming signal in advance, the entropy of the signal is low; if he is fully ignorant, the signal has a high entropy [3]. The goal of our ADC architecture is to use this pre-knowledge in the process of A/D conversion for saving energy as well as for data compression.

The pre-knowledge on the input signal is made available to the entropy-coding A/D converter ECOC by way of a probability distribution (PDF) and a prediction of the next input sample. The ECOC uses this information in order to determine the code for the next input sample efficiently [2]. The successive-approximation algorithm used has many similarities to Huffman coding, which is well-known from digital data compression [4]. In each step of the algorithm, the sampled input is compared to an optimal threshold that is determined on the basis of a Huffman code tree. Input values with a high probability only need a small number of comparisons, while more comparisons are necessary for input values with a low probability. On the average, the number of comparisons ideally becomes equal to the entropy. Since power consumption is mainly proportional to the number of comparisons, energy is saved by this algorithm because the converter returns to power-down earlier in the fixed sampling interval.

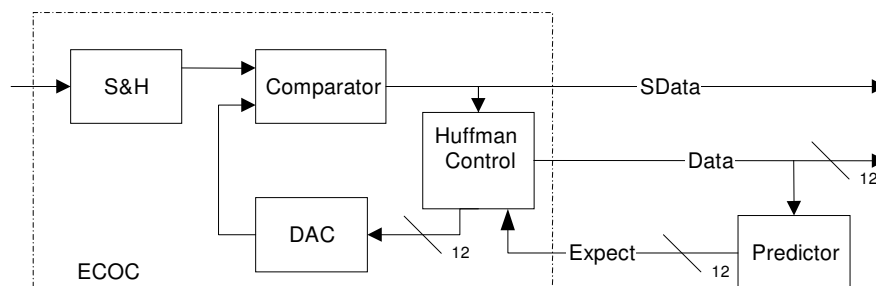


Figure 1: Block diagram of the entropy-coding A/D converter

The block schematic diagram of the ECOC (Figure 1) shows the components of the system architecture [2]. It consists of the blocks *Comparator*, digital/analog converter (*DAC*), sample&hold circuit (*S&H*) and *Huffman control*, which executes the quantization algorithm. The *Expect* input allows the control unit to take into account an estimation of the next input sample, which is e.g. provided by some prediction mechanism (an example is discussed in the next section). The parallel output *Data* delivers the (full binary) code corresponding to the analog input value. The single-bit serial output *SData* is connected to the comparator and yields the compressed Huffman code. It can be directly used to transmit the digitized input signal over band-limited channels like a radio link in a wireless sensor network.

Because of restrictions of chip size and manufacturing cost, the ECOC currently accepts only three-region piecewise-constant probability distribution functions as shown in Figure 2. The width of the high-probability region as well as its center can be configured externally. The implemented conversion algorithm consists of two phases. In the first phase, the high-probability region is processed with the conventional binary search strategy. Only in case that the input value is found to lie outside this region during the first phase, the full input range is searched in a second phase. For instance, if the entropy of the input is 6 bit, while 12 bits are needed for full resolution of the input range, a 12-bit ECOC can be configured to 6 bit for the high-probability region. Then, in most cases the input value is converted with only 6 comparison operations, while in the unfavorable case 18 (6+12) comparisons are needed. On average, however, the comparison effort is e.g. only 6.6 bits. By this simplification, there is no need




Figure 2: Piece-wise constant approximation of the probability distribution

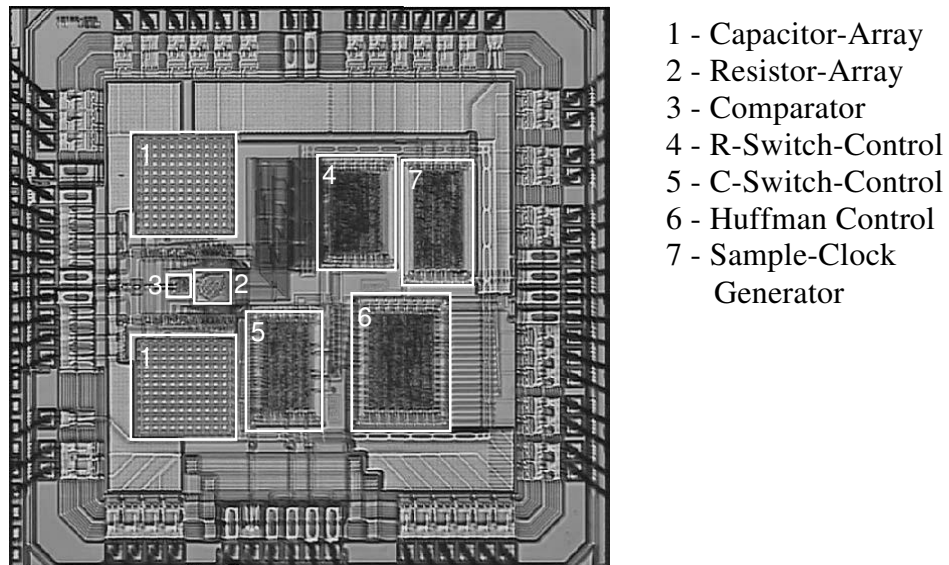


Figure 3: ECOC chip photograph

to store a large Huffman code tree [3],[4], which would be necessary for arbitrary probability distributions.

We realized a prototype implementation of the ADC with 12 bit resolution in a 0.6 μm CMOS technology [2]. A photograph of the realized chip is shown in Figure 3. The core area containing the ADC components has a size of 3 mm^2 , the size of the whole chip is 10 mm^2 .

Performance results measured on the prototype are shown in Table 1 [2]. A maximum sampling rate of 200 kS/s has been achieved, while the signal-to-noise-and-distortion-ratio (SINAD) was 66 dB. In the conventional binary-coding mode, power consumption of 1.4 mW was obtained. For an experimental investigation of the entropy-coding mode, we computed zero-mean Gaussian white noise waveforms of well defined entropies and loaded these into an Agilent 33120A Arbitrary Waveform Generator. This instrument generated a corresponding analog noise signal, which was applied to the input of the ECOC. The compression factor as well as the power con-

CMOS Technology	0.6 μm
Resolution	12 bit
Max. Sampling Rate	200 kS/s
SINAD	66 dB
Max. Compression (2 bit input)	4x
Power (binary coding)	1.4 mW
Power (entropy-coding, 3 bit input)	0.8 mW

Table 1: ECOC performance (all values for 200 kS/s sampling rate)

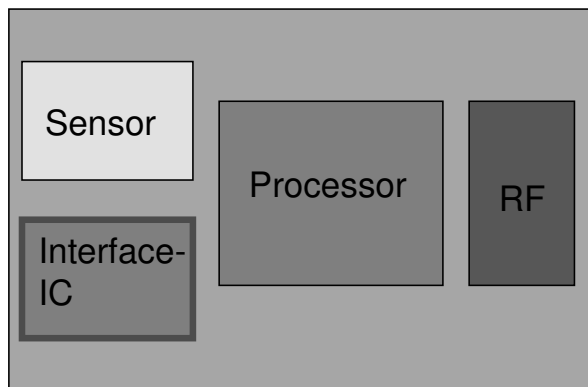


Figure 4: Sketch of wireless sensor node

sumption was measured. We found that the integrated data compression comes close to the Shannon limit down to one bit distance. The power consumption is quite low with $800 \mu\text{W}$ in the entropy-coding low-power mode at full conversion speed.

The adaptability, the low power consumption, and the integrated data compression make the converter very well suited for wireless sensor networks. A sketched view of a possible wireless sensor node is shown in Figure 4. The node consists of 4 main components, i.e. the sensor, an integrated circuit (IC) including the sensor interface circuitry (amplifier, filter etc.) and the ADC, a microprocessor, and a radio unit. While our current prototype-IC contains only the ECOC (see Fig. 3), it is useful for the purposes of miniaturization and power-saving to integrate the sensor interface and the ECOC together on a single application-specific circuit (Interface-IC in Fig. 4). Depending on the application, it might even make sense to integrate the microprocessor on the IC as well. The usefulness of the ECOC on wireless sensor nodes like this will be further illustrated by investigating a hypothetical application scenario in the next section.

3 Application Scenario: Distributed Wave Detection

The properties of the entropy-coding A/D converter can be utilized favorably in such wireless sensor network applications where the network maintains a global model of the sensor data, as described e.g. in [5]. In that paper, the authors report on their development of a distributed regression algorithm, which is executed on the sensor nodes to fit a global function to local measurements. Details on an evaluation of the algorithm based on data from a sensor network at the University of California, Berkeley are also given. The authors point out that using a global model of the sensor data, the communication workload can be drastically reduced, since only the model parameters instead of all the sensor values need to be transmitted to a central station. Raw sensor data have to be exchanged only locally between neighboring nodes during the execution of the distributed regression.

In the context of the present paper, a particularly important aspect of the concept described in [5] is that after the initial setup of the global model, each node can predict the behavior of the sensor field in its local area [5]. Thus, if entropy-coding ADCs are used in the sensor nodes, not only the communication workload is reduced but also the




Figure 5: Distributed wave detector

cost of A/D conversion. If the prediction from the global model is supplied to the *Expect* input of the ECOC (cf. Figure 1), the converter is enabled to acquire the subsequent local sensor values very efficiently.

Other advantages of the global sensor data model mentioned in [5] are supported by the ECOC as well. As a byproduct of its operation, the ECOC indicates automatically if the converted value is close to the prediction or not, because a dedicated output bit signals whether the value was found inside the high-probability region or not. This feature facilitates the *Adaptive Sampling* and the *Outlier Detection* mentioned in [5], where *Adaptive Sampling* refers to suppressing messaging of values that are near to the prediction. Outliers can be identified by the ECOC when the above-mentioned output bit indicates that the value was found outside the high-probability region, i.e. was unexpected. This information can be used to detect sensor failure or alarm conditions [5].

Let us elaborate on the concept of maintaining a global model of the sensor field a bit further by considering the following hypothetical application scenario (see Figure 5). The task is to monitor some area (or volume) of interest for the occurrence of a certain type of waves of the sensed quantity (e.g. pressure), in which case an alarm has to be triggered and the propagation of the waves has to be followed up precisely. It is assumed that in the normal state some noise of low amplitude is present in the pressure field while the waves of interest have higher amplitude.

For this purpose, a wireless sensor network is laid out in the area. Each sensor node contains a pressure sensor, sensor interface electronics, an entropy-coding A/D converter, a microprocessor, and a radio unit (cf. Fig. 4).

A distributed wave detection algorithm is implemented on the network, which might operate in a way as described in the following. Normally, all sensor nodes are put into

1-bit resolution. This means that the value from the pressure sensor is compared by the ECOC with a threshold only, which is defined above a suitably selected maximum expected noise amplitude. The nodes are "armed" (state *green*). Since they have to do only a single comparison per sample, this operation can be carried out with very low power consumption. Also, no radio transmission is necessary in this state.

If the sensor value is found to be above the threshold, it has to be detected if the wave is present or if the noise has just an unexpectedly high value. Remember that an alarm has to be triggered only if the wave is detected; higher values of noise should go unnoticed. The wave detection is performed by several sensors in the network in cooperation. Corresponding decentralized detection algorithms have been discussed for instance in [6], [7], and [8]. These papers describe how a nonlocal detection algorithm can be distributed and carried out by several nodes of a sensor network cooperatively. In the scenario considered above, after finding the sensor value greater than the threshold, the sensor node digitizes the value with higher precision and sends the result to a few neighboring nodes. By this, the local network neighborhood goes into a state of alertness (state *orange*), meaning that suspiciously high amplitudes have been found in the local region, and the actual distributed detection has to be carried out. The sensors in state *orange* now themselves convert their sensor readings with higher precision and send their values to their respective neighbors. (It might be advantageous to keep the precision and the sampling rate in state *orange* at a low level in order to not consume unnecessarily much power. The details depend on the actual implementation of the detection algorithm.)

Now that all nodes in the local region have knowledge of the values of their neighbors, they apply the decentralized detection algorithm to these values. This could be a Bayes detector testing the values against the hypothesis that a wave is present. During this test, the spatio-temporal correlations of a wave, which are not present in noise, are taken advantage of. If the test result is unfavorable (i.e. no wave present), this outcome is communicated to the neighbors, and if all nodes agree with this, the network goes back into state *green*. If not all nodes agree, the nodes stay in state *orange*, and the process is repeated.

If the test outcome is favorable (i.e. a wave has been found), the node goes into state *red*, converting the analog sensor value now with full precision and possibly higher sampling rate in order to observe the wave closely. It also estimates some parameters of interest of the wave (e.g. amplitude, frequency, wave length, direction of propagation, or location of the wavefront). All this information is communicated to the *orange* neighbors, which puts them into state *red*, too. The other sensors also estimate the wave parameters based on their local information and exchange the results with the neighbors. A distributed estimation of the wave parameters can be carried out now, using e.g. an algorithm similar to the one described in [5], or using a Bayes-based parameter estimation that integrates several sensor readings into a common model (multisensor data fusion) as discussed in [9].

The local Bayes estimators integrate all this information into estimations with increasingly higher confidence. As soon as a sufficient number of sensors in the *orange* neighborhood agree on the presence of the wave and its parameters, the *red* state is flushed over the whole network, and an alarm to the user of the network is triggered. From now on, all sensor nodes are informed on the parameters of the wave, and they can predict the coming wave with good confidence (it is assumed here of course that the speed of information transfer in the network is higher than the propagation veloc-

ity of the wave). The sensor nodes adapt their entropy-coding ADCs according to the predictions and use the mechanism described in Section 2 for monitoring the wave with precision but also with low power. In applications where the user is interested in all the local sensor readings during an alarm, the compressed output of the ECOCs can be transmitted to the central station very efficiently. Otherwise, only the parameters of the wave model have to be sent.

During the alarm, the wave parameters are continuously updated. If the wave parameters indicate disappearance of the wave, the sensor network returns back to state *green*.

The advantages of the outlined distributed wave detection compared to local or central detection are

- sensor values are transmitted only in the neighborhood,
- computation takes place mostly local,
- distributed detection is based on spatio-temporal correlations of the wave and thus more reliable than local detection (less false alarms),
- only global information as the wave parameters are broadcast to the network,
- the disseminated information on the detected wave is used to monitor the wave efficiently.

The entropy-coding A/D converter described in Section 2 is particularly suited to this type of scenario. In state *green*, the ECOCs might be put into power-down mode most of the time. Their configurable resolution allows to adapt the number of processed bits to the current state of the node, reserving high resolution and thus high power only to the case of an alarm. As soon as the wave parameters are determined and broadcast, the spatio-temporal waveform at each moment and at all node positions can be predicted very well. The high-probability regions of the ECOCs then can be adapted accordingly, thereby saving power even in the case of an alarm. Higher power is spent only for capturing additional information on the wave insofar as the actual wave deviates from the predictions. Raw sensor readings are delivered by the ECOCs also in compressed format, which keeps the cost small of transmitting the values to the neighbors or to the user.

While it is difficult to determine the efficiency without the actual development and investigation of the system for a specific application, a rough indication on the energy-saving potential for this scenario can be given as follows. For simplicity, let us assume a constant sampling rate of 20 kS/s, a basic entropy-independent power consumption of 15 μW , and an entropy-dependent power consumption of 8 $\mu\text{W}/\text{bit}$. These numbers can be predicted for the next ECOC generation from the previous experiments. A conventional continuous 12 bit data acquisition would then consume 111 μW per ADC. If we further assume that the sensor nodes work in full 12 bit for 0.1% of the time (in state *orange* and in the beginning of state *red*), in 3 bit for 1% of the time (state *red*), and in 1-bit mode for the rest of the time (state *green*), the ECOC can be estimated to consume 23 μW on the average, which amounts to an energy saving of 79%. This factor includes only the power reduction of the A/D conversion itself; further savings are possible by communicating only the global or the compressed data instead of the full sensor readings.

4 Conclusion

An entropy-coding analog/digital converter has been presented that combines A/D conversion and data compression in a single process. The converter utilizes predictions or other prior knowledge of the input signal for achieving low-power operation. Measurements on a prototype CMOS chip have been reported.

The converter can be tailored to specific applications and integrated together with sensor interface electronics and a microprocessor on a single chip for very small-size and low-power sensor nodes. It was shown that the concept of entropy-coding A/D conversion can be utilized most advantageously on sensor networks that maintain a global model of the sensor data field. This was illustrated further discussing a scenario of distributed wave detection.

Currently we are working on the next ECOC generation. This version will have externally configurable resolution between 1 and 12 bits. Further, the so far relatively high standby power will be reduced, which will bring forth a power reduction factor close to the data compression factor. The new version will be very flexible and can be adapted in terms of resolution, conversion rate and input entropy to various applications. Particularly in the entropy-coding mode it will consume extremely small power. All configurations, including a power-down mode, can be changed online during the operation. This feature allows to adaptively react to changes of the situation and to put the converter in each case into the optimum configuration.

References

- [1] K. Uyttenhove and M. Steyaert. Speed-power trade-off in high-speed ADC's: What about nano-electronics? In *Custom Integrated Circuits Conference (CICC 2001)*, pages 341–344, San Diego, California, May 6-9, 2001.
- [2] R. Peck and D. Schroeder. A low-power entropy-coding analog/digital converter with integrated data compression. In *Proceedings of the 29th European Solid-State Circuits Conference (ESSCIRC 2003)*, pages 173–176, Estoril, Portugal, September 16-18, 2003.
- [3] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
- [4] A. Gersho and R.M. Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston, 1992.
- [5] C. Guestrin, P. Bodik, R. Thibaux, M. Paskin and S. Madden. Distributed regression: an efficient framework for modeling sensor network data. In *IPSN'04*, Berkeley, California, April 26-27, 2004.
- [6] R.R. Tenney and N.R. Sandell. Detection with distributed sensors. *IEEE Transactions on Aerospace and Electronic Systems*, AES-17, pages 501–509, June 1981.
- [7] H.J. Kushner and A. Pacut. A simulation study of a decentralized detection problem. *IEEE Transactions on Automatic Control*, AC-27, pages 1116–1119, October 1982.
- [8] Z. Chair and P.K. Varshney. Optimal data fusion in multiple sensor detection systems. *IEEE Transactions on Aerospace and Electronic Systems*, AES-22, pages 98–101, January 1986.
- [9] C.R. Smith. A Bayesian approach to multisensor data fusion. *Proc. SPIE*, vol. 1699, pages 285–299, 1992.