

# Modelling and simulation techniques for highly integrated, low-power wireless sensor networks

B. Otis, S. Gambini, R. Shah, D. Steingart, E. Quévy, J. Rabaey, A. Sangiovanni-Vincentelli and P. Wright

**Abstract:** The design of state-of-the-art low-power wireless sensor nodes involves the convergence of many technologies and disciplines. Submicron complementary metal oxide semiconductor (CMOS) devices, micro-electro-mechanical system filters, on- and off-chip electromagnetic elements, sensors and thin-film batteries are some of the technologies that will enable pervasive systems such as wireless sensor networks. High system complexity requires the use of many simulation environment during design: algorithm simulators, mechanical finite element analysis, behavioural and transistor-level circuit simulators, electromagnetic (EM) simulators, thin-film battery simulators and network simulators. It is shown that highly integrated, self-contained systems require multiple-domain simulations to uncover complex interactions between domains. Specific examples of block- and system-level design methodologies used in low-power wireless systems are presented here. Bottlenecks in the current methodology will be identified with an eye towards improving the scope and resolution of system-level simulations.

## 1 Introduction

Recently, the emerging concept of Ambient Intelligence has promised to change the way human interact with technology [1]. This technology will be reactive and distributed, affording high levels of redundancy and fine spatial granularity. The heart of this movement is rooted in extremely small, inexpensive and ubiquitous wireless sensors. The level of difficulty of implementing a suitable hardware platform is severe. Each sensing element, or node, must be physically small (approximately  $1 \text{ cm}^3$ ), inexpensive (to allow high network density) and completely wireless and self-contained (to minimise installation and maintenance costs). Therefore the nodes must operate for their entire lifetime ( $>10$  years) on a primary battery or energy harvested from the environment. The difficulty of this requirement is exacerbated by the lengthy list of required node functionality: each node must contain ample digital computational power, environmental sensing, a highly integrated power train, a wireless transceiver and sufficient intelligence to form reliable *ad hoc* networks [2].

To achieve this level of functionality, cost and miniaturisation, all components should be batch fabricated using thin-film processes. Creative and robust packaging is necessary to allow electrical connectivity between chips and physical connectivity to the outside world. See Fig. 1 for an example of a  $1 \text{ cm}^3$  sensor node.

To allow diverse configurations in the sensors and power sources of the nodes, a modular approach is desirable. This would allow, for example, a solar energy scavenger for outdoor applications and a vibrational energy scavenger for dark but noisy industrial environment. This is a substantial departure from traditional chip-on-PCB packaging. In addition, each of the technologies that are integrated on this platform is experimental and unproven in isolation. This increases the level of uncertainty and the need for comprehensive modelling and simulation across the various subsystems.


Another difficulty is the multiple-domain nature of this design problem. Techniques that were used to model and simulate this complex system are described. First, challenges faced during the modelling, simulation and design of the power train of the sensor node will be described.

## 2 Power distribution


The heart of a sensor node is the power train. This system must harvest energy from the ambient environment and converts it to useful electrical energy to power the electronics. The efficiency, form-factor and drive capability of the power train are of great importance. There are three main components that comprise the power train (Fig. 2):

1. Energy generation – An energy harvester, primary (non-cycling) electrochemical battery or reactor.
2. Energy storage – A secondary (cyclable) electrochemical battery or super capacitor.
3. Power conditioning – The circuitry necessary to transform the energy generated to useable power to be stored or to run the device. This is also the circuitry responsible for storing ‘excess energy’ from the generator.

Power train design is particularly difficult for sensor nodes because of the power, size and environmental constraints of the node. Let us consider as a working example a total device volume of  $1 \text{ cm}^3$ . Given a short term and/or extremely low duty-cycle application (e.g. 1 year lifetime



**Fig. 1** Conceptual drawing of  $1 \text{ cm}^3$  sensor node



**Fig. 2** Power train overview

of sending and sensing once an hour), all three components listed above could be met by a single primary battery. However, if the node must sustain higher duty-cycles and last for upwards of 10 years, the design becomes much more difficult. Given the latest in wireless sensor network technology, no primary battery exists that can meet this goal in  $1 \text{ cm}^3$ . The three components listed above must be divided physically into three units, which must then be volume optimised for the given application. The simulation and modelling of this multiple-domain system is crucial for this optimisation. This is particularly difficult because this analysis must capture the physical (vibrational scavenger), photonic (solar cell), chemical (thin-film battery) and electrical (power electronics) domains accurately.

## 2.1 Electrochemical energy storage

This section provides a brief introduction to state-of-the-art thin-film energy storage technologies. A discussion of other energy storage and scavenging technologies can be found in [3].

**2.1.1 Primary batteries:** Outside of a nuclear reaction, primary electrochemical cells provide the greatest energy density available. An electrochemical energy cell consists of two complimentary redox reactions in electrodes separated, chemically, by a potential gradient, and physically by a layer referred to as an electrolyte. When a load is placed between the terminals of a cell the circuit is closed, the electrodes become mutually aware of the others' presence due to the potential gradient. Thermodynamics demands that this potential gradient diminishes and an ion–electron pair is created depending on the chemistry, the ion electron pair may come from a single electrode or from opposite electrodes. The electrolyte allows only the ion to pass, so the electron must flow through the load, thus powering the circuit. The chemistry of the cell determines the potential change during discharge as well as the potential, energy density and power density.

**2.1.2 Secondary batteries:** Secondary electrochemical batteries are batteries that can be cycled multiple times. Cells that are managed well with carefully chosen chemistries can last well over 1000 cycles. The modelling of this cycling behaviour is very difficult, only now are ab initio models coming to light that can adequately predict cycled fade; empirically determined estimations and extensive testing are still the typical methodology for determining cycle life.

**2.1.3 Electrochemical (super) capacitors:** Electrochemical capacitors, or super capacitors, transport charge internally as batteries do, via ionic conduction through a salt electrolyte, but store charge as a capacitor, by maintaining an adsorbed charge on a surface. Ions pack more readily than electrons; therefore these capacitors can hold more charge per unit area than a regular capacitor. However, since these ions must travel through a dense liquid to provide power, super capacitors cannot discharge as quickly. A major benefit of super capacitors over batteries is that they can cycle hundreds of thousands of times at relatively high discharge rates.

## 2.2 Battery modelling and simulation

This section describes the state-of-the-art and existing challenges in modelling and simulating thin-film batteries.

**2.2.1 Battery capacity:** Battery modelling is a difficult task, but can be broken into bounded tiers to provide adequate general expectations before extensive testing is required. A brief overview on the qualitative aspects of cycling will now be given. The absolute theoretical capacity of a cell is determined by the amount of active material within a cell. The energy density of a cell is limited by the amount of active material a single electrode can hold. Power delivery is far more complex. At an atomic level, the limiting factor is the speed at which an ion–electron pair can be generated. Cell performance is determined by the thickness of material the ion must travel through, the temperature of the cell and the electrical load. Ohmic losses reduce the battery potential as loads increase. In lithium ion cells, the electrodes are often poor electronic conductors; thus a trade-off exists between designing for maximum capacity by making a thicker electrode and designing for minimum ohmic losses and making a thinner electrode. Engineering solutions exist to ease this trade-off, but modelling exactly how these electrodes interact is difficult. Luckily, software that does this accurately has been developed, and has been consistently improving over the last decade. Finally, accurately predicting cycle fade, from a first principles approach, is extremely difficult as it involves couplings between the electrical, chemical and mechanical domains.

**2.2.2 Empirical models:** Some battery types are so common (e.g. AA, 9 V) that it is worthwhile to run many cells through many different loads and cycle tests to create curve fit equations to describe their behaviour. This approach has been discussed extensively. The problem with such a black box approach is that the resulting model is hardly parametrisable and therefore does not allow the designer of on-chip thin-film batteries to predict at design time the effect of altering component geometry or reaction chemistry.

**2.2.3 Phenomenological models:** The standard model for lithium ion batteries is DUALFOIL, a FORTRAN

program developed by John Newman at UC Berkeley. Although computationally heavy, this software package allows for one-dimensional simulation of a cell with control of over 60 parameters. It is possible to simulate a cell with an incremental change in any of those factors, thus with a few hours of script writing and a couple of days of computer time, a significant usage matrix can be compiled for a battery. This model excels at predicting charge/discharge times based applied currents or potentials. Cycle fade modelling can be added on top of as described previously. The following section shows a practical application of DUALFOIL.

**2.2.4 Simulation results:** We now see how some of the aforementioned accurate battery models enable evaluation of the interaction between the power train and the communicating and sensing elements of the sensor node, enabling global lifetime optimisation.

Consider a sense and send wireless sensor node has a choice between three radios having a maximum transmit current of 8, 12 and 22 mA, respectively, at a nominal potential of 2.5 V. Transmission time is estimated to be 50 ms. The node has a sleep current of 80  $\mu$ A, and the sensing draws an average current of 5 mA for 20 ms. The battery for the sensor is to be no larger than 400  $\mu$ m by 1  $\text{cm}^2$ , and should use a gel-type (pseudo-solid) electrolyte.

Fig. 3 shows a simulation of the cell for all three radios for sense and send periods ranging from 15 s of sleep to 1 h of sleep. The results predict that even, given the relatively high sleep time of the cell, decreasing the peak load has a large effect on total run time, particularly at short sense and send periods. The simulation also shows the considerable effect of ohmic losses in a cell. Fig. 4 shows the energy delivered against sense and send period for all three. The use of the 22 mA radio actually causes a decrease in total delivered energy, even at longer sense and send periods. This energy loss per cycle is also qualitatively indicative of cycle fade. Thus, for efficient operation, a low maximum current draw is preferable to bursty, high-current operation given a constant amount of energy dissipation.

To further illustrate the effect of ohmic losses in a cell, let us compare the pulsed current draw with an average current draw that is equivalent to the total coulombs passed (Fig. 5); this load levelling can be accomplished by capacitive filtering. When slowly discharged, all three cases realise a maximum capacity of 120 W-h/kg. When pulsed, the 22 mA radio never achieves this utilisation, even at hour-long sense




Fig. 3 Discharge times for various radios

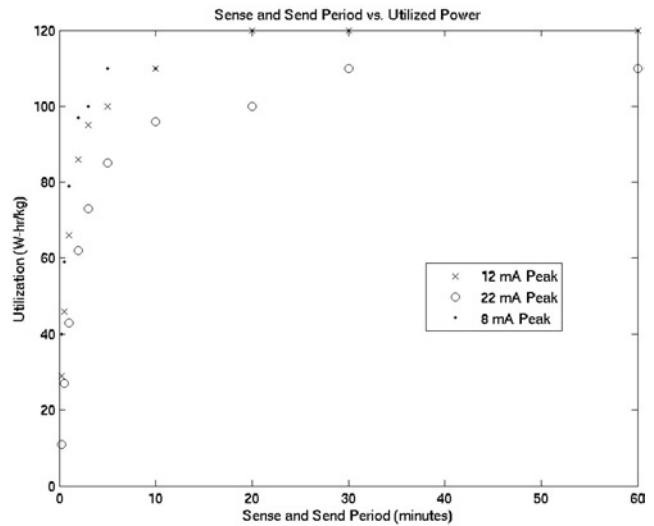


Fig. 4 Battery utilisation for various radios

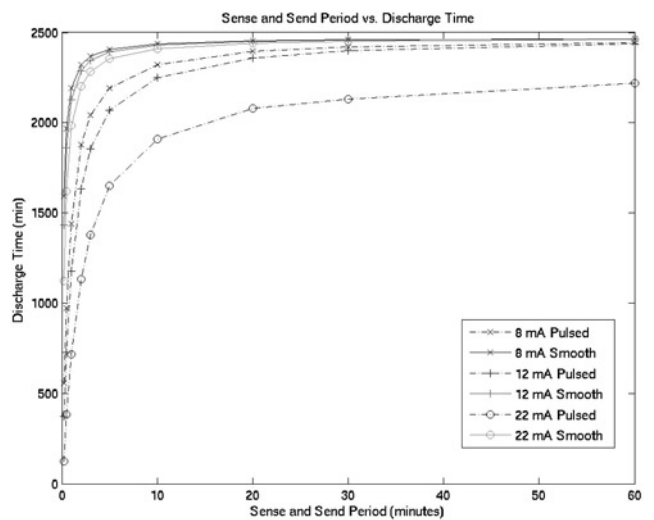


Fig. 5 Discharge time for direct pulsed and load levelled systems

and send intervals. We see that the 12 mA radio reaches maximum utilisation at 20 min per sense and send.

This model can be further extended to account for charging efficiencies by modelling expected currents from the selected method of energy scavenging. However, there is currently no tie-in between the thin-film battery simulator and a transistor-level circuit simulator. This connection would allow the evaluation of circuit performance over battery charge degradation.

### 3 Communication and network algorithms

Network performance is the ultimate metric for a system, since it encompasses all the other subsystems and components: individual protocols, radio channels, analog and digital radio performance and other processing/memory elements. This also makes it a challenging problem, since the optimisation space is enormous and to make the problem tractable, many abstractions are typically made.

There are mainly three approaches taken to identify overall network and communication algorithm performance. In increasing order of accuracy, these are:

1. Analysis,
2. Simulation,
3. Testbed implementation.

As can be expected, there is a direct trade-off between accuracy and the amount of effort required. However, all three approaches are necessary to understand the performance of the system as a whole. Analysis provides an intuitive idea of system performance, but it is typically limited to simple cases with idealised assumptions. For example, performance of medium access (MAC) protocols can be analysed, but using ideal assumptions such as Poisson traffic arrival and uniform random distribution of nodes. This approach provides information of performance within a single neighbourhood (a radio range of a node), and is usually limited to average case behaviour such as average throughput and delay.

On the other hand, simulation offers the opportunity to look at the performance of networks of nodes (in the order of hundreds of nodes). However, it is impossible to optimise each parameter of the system to maximise performance, so many parameters are chosen based on analysis or individual component simulation done elsewhere. Additionally, detailed simulations are time intensive, so suitable abstraction choices are made. For example, for network simulations, only simple bit-error or packet-error models are used, rather than detailed radio models or hardware models of the individual nodes. Obviously, the choice of models is critical in ensuring the validity of the simulation results.

Once a rough estimate of the system performance is obtained, the network is usually tested in a realistic testbed environment. This involves configuring the actual hardware and software, and monitoring the system performance. Although the performance results obtained in this fashion are actual results, it is very difficult to monitor every node in the network, and to correlate the results across nodes due to the copious amounts of data produced. Additionally, it is often impossible to change many parameters of the system, so apart from some software tuning, implementation can only be used for system validation.

For the PicoRadio system, an extremely simple network level analysis in [4] was used, which provided the average throughput and delay achievable gave insight into basic system behaviour. Network simulation [5] was then performed in Omnet++ [6] using packet level error models detailed in [7]. This error model abstracted away the actual performance of the radio, as well as link level details such as forward error correction and cyclic redundancy check (CRC). Using the simplified packet error model enabled network level simulations of a large number of nodes with all the complex MAC layer and network layer interactions among them. Having validated the protocols using simulation, they were then implemented on a testbed consisting of hardware nodes to obtain their performance in realistic scenarios. It would greatly benefit the accuracy of the results and the insight provided by them to incorporate more complexity into the network simulations, such as realistic transceiver physical layer models. The trade-off between evaluation effort and accuracy is shown in Fig. 6.

#### 4 MEMS/circuitry co-design

With several commercial successes and continued emergence of new products and applications after more than 25 years of R&D, micro-electro-mechanical systems (MEMS) encompass a wide range of engineering specialties, such as optical-mechanical actuator design for displays, micro-fluidic for lab-on-chip analysis and signal processing. The difficulty in assembling standardised, widely available design flows is exacerbated by the

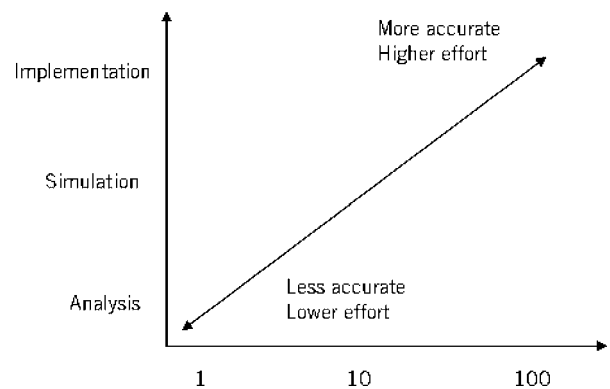


Fig. 6 Effort, accuracy and network size trade-offs in network performance evaluation

multitude of specialties and competences required, such as mechanical design, process technology engineering and broader general physics. In the case of wireless sensor network applications where size and power are the primary concerns, MEMS resonators present the compelling advantage of a very small form-factor along with the ability to be monolithically integrated with CMOS electronics. Over the past few years, high-Q micromechanical resonators of a variety of mechanical designs and transduction mechanisms have been demonstrated [8–11]. Although featuring performances similar or superior to existing off-chip technologies such as quartz and ceramics resonators, these devices promise to implement filtering functions in communication systems with substantial savings in power, system volume and cost. Overall, functions where MEMS resonators can allow for efficient and complete integration are radio frequency (RF) channel select filters, RF carrier generation and frequency references. The designer's scope for implementation of those devices consists of process technology, electro-mechanical design and system-level design.

#### 4.1 MEMS design flow: from multi-domain modelling to library-based parametric design

The initial step in implementing MEMS resonators within a system is considering the system-level performances to be achieved by the mechanical structure, whether a crystal [12], a filter or a mixer [13]. Fig. 7 shows the design flow usually experienced at the early-stage definition of the resonator topology.

Considering the system-level specifications, an experienced designer will choose a particular topology and transducer configuration available by the technology, and will then rely on simple analytical modelling to iteratively funnel the choices in terms of shape and dimensions, and

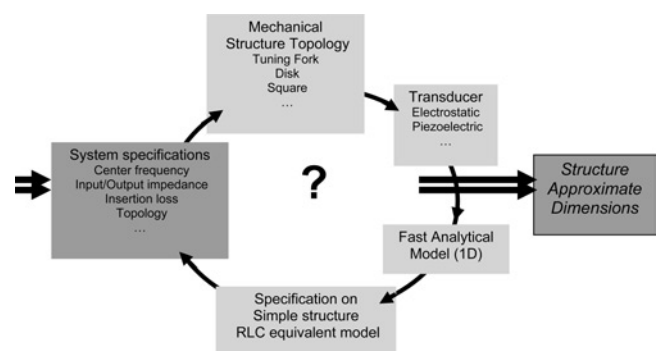
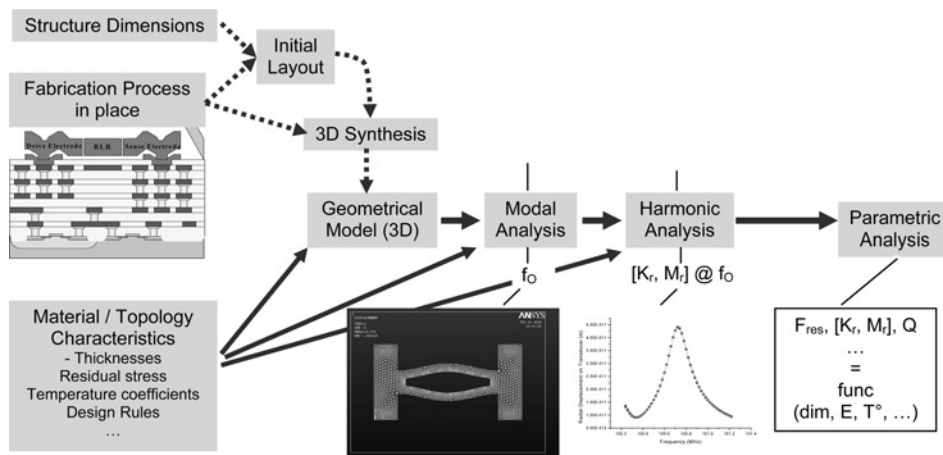



Fig. 7 Preliminary specification-based MEMS design flow



**Fig. 8** Electro-mechanical finite-element modelling and optimisation

extract performances matching the system requirement to the first order of the target application. Once the basic topology (shape, transducer, differential/single-ended ports) and coarse dimensions have been defined, extensive multi-domain (electromagnetic, mechanical) finite-element analysis of the structure is undertaken to generate an accurate model of the device, as shown in Fig. 8. Combining knowledge about the fabrication process flow, structure configuration and material properties, finite-element simulation will help extract equivalent electrical and mechanical parameters under certain conditions to verify their compliance with system specifications. Specialised modules and parametric analysis are also usually used to optimise and model more accurately certain parameters, for example acoustic losses in the device [14], and fit them with experimental data.

To the designer, this represents a rather complex flow for the optimisation of such a structure within the target system, as it involves the simultaneous work of specialists in technology, material and mechanical design – especially if the MEMS device is to be re-invented for any new system. It is thus necessary for the electrical engineer to narrow the design space at the system level by limiting the number of structure topologies available in a given technology, while keeping flexibility through parametric design. This approach is largely inspired from the approach taken for transistor device design where only a few constrained structures are available for circuit design. Conveniently, MEMS structures and resonators in particular can be modelled rather accurately by equivalent lumped element resistive, inductive, and capacitive (RLC) network, under certain assumptions like linearity of the constitutive physical equations. As shown in Fig. 9, from the extraction of lumped mechanical characteristic such as stiffness ( $k_r$ ), mass ( $m_r$ ), quality factor ( $Q$ ) and transducer efficiency ( $\eta$ ), an electrical equivalent circuit is built, including  $R, L, C$  parameters depending on geometrical




**Fig. 9** Electro-mechanical resonator equivalent RLC circuit

dimensions, temperature, material properties and bias conditions.

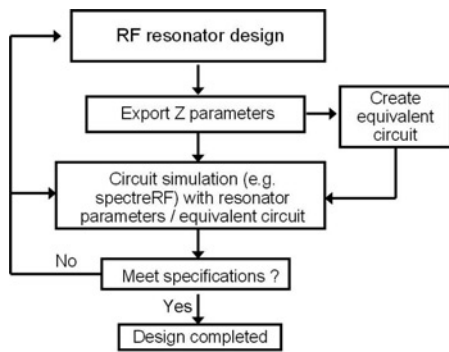
Such a model can be made more sophisticated by the use of behavioural equations coded in analog - hardware description language (HDL) to take into account nonlinearity or any other behaviour extracted from experimental device characterisation, such as noise. Then, the designer can rely solely on a library of MEMS structures that have been extensively modelled and characterised initially at least once, and implement them along with transistors or other integrated passives like inductors. The same way circuit design takes advantage of SPICE models for transistors, parametric modelling of MEMS structures allows performance optimisation by co-designing and optimising circuitry and MEMS resonators at the same time.

#### 4.2 RF MEMS/CMOS co-design

There is great interest in designing MEMS filters and switches for use in communication systems at very high frequencies. These RF MEMS devices can potentially decrease the power consumption and form-factor of RF transceivers by processing signals using high-quality mechanical devices instead of power-hungry electrical components. In addition, since the devices are batch-fabricated on semiconductor substrates, they are small and potentially integratable with the electronics. One example of this technology is the bulk acoustic wave (BAW) resonator, which provides a very



**Fig. 10** Bulk acoustic wave resonator photograph and measured impedance



**Fig. 11** Design flow for RF MEMS/CMOS co-design

high-quality factor resonance ( $>1000$ ) at RF frequencies [8]. Fig. 10 shows the measured impedance of a BAW resonator. The inset shows a photograph of a single resonator.

This resonator presents electrical impedance that varies three orders of magnitude over a few percent of its bandwidth. This response, which would be very difficult to produce using electrical components, provides high-rejection filtering at very high frequencies. One application of this technology is in low-power RF oscillators. The high-quality factor allows a very low phase noise for small bias currents. Additionally, since the resonators are batch fabricated and defined lithographically, the designer can define multiple sizes and conductor layouts on each chip.

These observations have led to a co-design of a BAW resonator and CMOS circuitry to develop a low-power, low phase-noise oscillator [15]. The goal of the co-design was to optimally size the mechanical and electrical device sizes to provide an optimised performance. The MEMS process used was a commercial, mature process. As such, the design process is similar to that of an integrated circuit process. First, a rough parameterised model of the resonator was used to provide an electrical circuit equivalent. Then, an electrical circuit equivalent was incorporated into the design equations of the CMOS oscillator. At that point, the analysis was verified using an RF circuit simulator. This design flow is summarised in Fig. 11.

Note that, in this design flow, the linearised electrical resonator equivalent discards substantial information about the mechanical resonator dynamics. For example, if the resonator were driven at high-power levels, a more complex electrical model would be required to capture any nonlinear behaviour. The optimisation variable of this design was power consumption. An optimised resonator and CMOS chip were fabricated, assembled and tested. A common figure-of-merit allowing a comparison of the phase-noise/power consumption trade-off shows this oscillator to be approximately two orders of magnitude more efficient than oscillators composed of on-chip LC tanks [16] due to the high-loaded quality factor of the electro-mechanical resonant tank.

### 4.3 Example: super-regenerative transceiver

The previous section described how RF MEMS technologies can greatly decrease the power consumption and reduce the phase-noise of high-frequency oscillators while providing a relatively controlled RF frequency reference. These observations have led to the development of a low-power RF transceiver using BAW resonators to define the carrier frequency on the transmitter and receiver (Fig. 12) [17].

A super-regenerative architecture was chosen, as it requires an RF oscillator in the receiver with a well-defined RF frequency. This requirement is particularly well suited to the benefits of BAW resonators. The design methodology of Fig. 11 was applied to the super-regenerative transceiver.

One substantial design challenge is the simulation of RF MEMS-assisted circuitry. First, this transceiver used BAW resonators with a high-quality factor and a high resonant frequency. Because of this, the duration of the transient start-up of the oscillator is long in relation to the period of the carrier frequency. In this receiver, the magnitude of the incoming signal modulates the start-up time of the RF oscillator. Thus, accurate simulation of the oscillator start-up time in the presence of noise and various RF inputs is crucial for an accurate performance prediction. Additionally, the operation of a super-regenerative receiver requires the RF oscillator to start-up multiple times for each received bit of information. Simulating the receiver operation over a number of received bits is computationally difficult. Simulation of complex RF oscillators using computationally efficient macromodels is an active area of research [18], which will help overcome this problem. Currently, the design methodology involves simulation to verify the performance of individual circuit blocks and behavioural analytical models to predict complex receiver performance metrics such as sensitivity.

## 5 Behavioural modelling of an ultra-low-power $\Sigma\Delta$ converter


Modern integrated systems demand increased performance and complexity of analog circuit blocks. This challenge calls for advances in modelling techniques as well as a robust design methodology. Analog design methodologies are traditionally *ad hoc* and rely on tremendous amounts of designer intuition and experience. The approach described here is more systematic and compliant with a platform-based design methodology [19], which is promising in decreasing design time and improving circuit/system optimisation. The solution of a design problem is found by first proceeding to a top-down architectural exploration approach, followed by a bottom-up component characteristic extraction and verification. The top-down phase is performed using approximate models, and allows quick examination of a variety of architectural alternatives. The bottom-up phase verifies that a certain set of properties holds for a component of the design.

### 5.1 Models for analog circuit design

Three models are used for the design of the  $\Sigma\Delta$  converter. The first (1) is an approximate model, used for high-level performance estimation and architecture selection. This model provides an entry point for the subsequent stages of circuit design and verification, reducing effort spent in these more time-consuming steps. This is referred to as ‘top-down’ design space exploration.

The second (2) component presented is an accurate behavioural model, used to reduce verification time of the designed system. This is the ‘bottom-up’ verification phase. The synergic fashion in which these two components should interact in any high-performance, high-complexity design is well illustrated in [19, 20] and summarised graphically in Fig. 13.

The third (3) is a behavioural model that allows relatively accurate prediction of the effect of various non-idealities one at a time. This can be a significant aid in the course of debugging unexpected measurement/simulation results



**Fig. 12** Super-regenerative CMOS/MEMS transceiver

that would be too time-consuming to perform with a transistor-level simulation.

The role of models (1) and (2) are shown in the first presented design, in which an approximate top-down model is used in conjunction with an accurate bottom-up model to perform high-level design-space exploration as well as rapid verification of a high resolution A/D converter. Strong and weak points of the proposed approach are outlined together with the resulting simulated performance. The second presented case study demonstrates the need of model (3), which is used to understand the mechanisms of nonlinear distortion cancellation in a high-speed continuous-time  $\Sigma\Delta$  modulator.

## 5.2 Design methodology

The requirements placed on the A/D converter in a wireless sensor network receiver are stringent in terms of extremely

low-power dissipation ( $<20 \mu\text{W}$ ) and operation from a down-scaled supply ( $<1 \text{ V}$ ). Although baseband filtering allows relaxing ADC requirements to 8 bits accuracy, the filtering operation is relatively power hungry, and limits the achievable data rate of the radio as described in [17]. A higher dynamic range A/D (12 bits) sampling at high rate would enable improved system performance. By employing a switched capacitor  $\Sigma\Delta$  modulator, quantisation noise is suppressed only in a small signal bandwidth, whereas interferers leak through the high-quantisation noise regions and can be eliminated digitally. The goal of this design was to push the boundary of low-power modulators and overcome the difficulties of low-supply voltage.

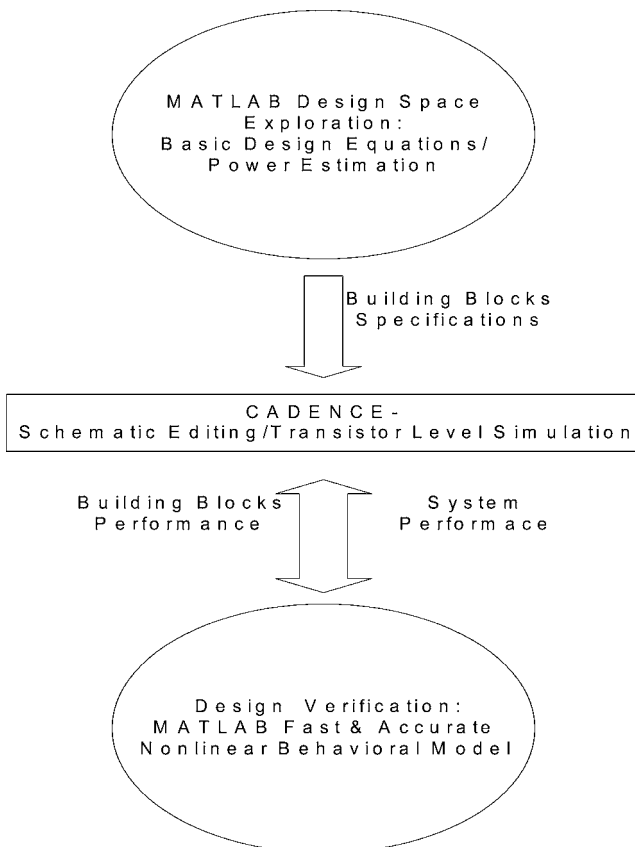
To achieve this goal, architectural decisions such as choice of modulator order, oversampling ratio, loop filter coefficients and forward quantiser resolution have been made in a power-aware fashion with the assistance of an *ad hoc* power estimation framework. This framework is described by the aforementioned model (1).

Although a correct architectural choice is important, true power optimisation also demands considerable effort at the circuit design level. The use of low-gain operational amplifiers was explored in this project as a way to minimise power consumption. Consequently, distortion performance of the converter is impacted in a measure that needs to be quantified accurately and quickly. This step requires the use of model (2) (a fast and accurate bottom-up verification model).

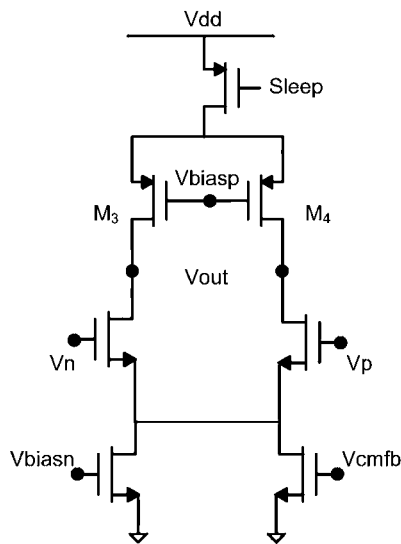
## 5.3 Power estimation

The power estimation tool comprises a custom graphical interface, used to input design parameters such as signal bandwidth, modulator order and desired closed loop resolution, as well as loop filter coefficients. This interface acts as a wrapper to a various MATLAB functions, which perform the core calculations.

The goal of this tool is to provide a realistic performance estimate without going through the process of designing the whole circuit in detail. This is one instance of the performance feasibility problem, one of the problems at the heart of analog computer-aided design (CAD). In this design, an approximate solution can be found by making a set of assumptions that is to be checked with a transistor-level simulation. In this case, the basic assumptions are that the transition frequency of the operational transconductance amplifier (OTA) input device is much higher than the sampling frequency  $F_s$ , so that the input differential pair can be biased at a low inversion coefficient  $<0.1$  [21],



**Fig. 13** Analog design flow for a low-power  $\Sigma\Delta$  converter



OTA Gain(dB)	40
GBW <sub>OTA</sub>	60Mrad/s
SlewRate	9.5 MV/s
Flicker Noise Corner $f_k$	15kHz
Noise Factor F	2
Output Swing	300mV

Fig. 14 OTA schematic and specifications

enabling operation at a high transconductance efficiency ( $G_m/I_d$ ) ratio of 25, and that capacitive load of each operational amplifier is dominated by its own sampling capacitance.

#### 5.4 Circuit design

From the power estimation phase, it is shown that 12 bits of effective resolution can be achieved in a 50 kHz bandwidth with a total power consumption of 15  $\mu$ W when using a second-order modulator with single-bit quantiser operating at an oversampling ratio of 128. The OTA simplified schematic and rough specifications (as output by the power estimation tool) are summarised in Fig. 14.

These specifications may be met by using the simple differential pair with active load shown above. Such a simple structure has the important benefit of having a non-dominant pole located at the input devices' non-quasistatic characteristic frequency  $\omega_{nqs} = 5 \omega_T$ . Since  $\omega_T > f_s/2$  ( $f_s$  is the sampling frequency), the amplifier is automatically stable with a large phase margin and exhibits an almost ideal single pole response. Low-power and low-noise performances are in this case traded for open loop gain, which is limited to about 40 dB.

Simulations show that specifications can be met with a very low bias current of 6  $\mu$ A. To further reduce power dissipation as well as area, an OTA time-sharing technique [22] is used, where a single amplifier serves as the active device for both of the loop filter integrators in different clock phases.

#### 5.5 Bottom-up verification

The verification environment has been developed to satisfy two basic requirements:

1. Simulation speed: verification times of the order of 1 week are unacceptable. Performing successful design without adding too much calibration circuitry or significant margins requires verification times of the order of a few minutes for a complete system.
2. Simulation accuracy of nonlinear behaviour: gain is the main limitation in low-supply designs. Linear analysis predicts operational amplifier gain as low as 30 dB can be successfully used in  $\Sigma\Delta$  modulators at the 12b level. However, lowering the gain of the OTA will increase distortion.

Therefore low-power design demands accurate modelling of nonlinearity arising from finite gain.

The effects of finite amplifier gain and finite settling speed are incomplete charge transfer from the sampling capacitor to the integration capacitor and non-zero charge leakage from the integration capacitor under non-zero input. These effects have been modelled in this work in a way similar to that reported in [23]. Nonlinear gain makes the efficiency in the charge transfer and the amount of charge leakage signal-dependent, so that no closed form expression is possible for the integrator response. A numeric solution is instead computed by assuming an integrator large signal gain model  $G(V) = V_{out}/V_{in}$  and using fixed-point techniques. Here, a piecewise linear model [24] for the integrator gain as a function of the input voltage is used. Time and frequency domain accuracy evaluation was performed using both MATLAB and a transistor-level simulation. Accuracy better than 1.5% is obtained on the large signal time domain sinusoid, translating into a 3 dB error in the overall SFDR prediction. When applied to a whole modulator, this model gives the results summarised in Fig. 15.

The spectrum predicted by the behavioural model is compared to a transistor-level SPECTRE simulation. A large discrepancy is observed in the DC noise floors, which results in vastly different predicted signal to quantisation




Fig. 15 Simulated modulator output spectra

noise ratio (SQNR) values. The poor low-frequency quantisation noise rejection simulated in SPECTRE is a consequence of charge transfer through the parasitic input capacitance of the time-shared OTA. This second-order effect appears as a positive feedback loop around the loop filter, which decreases the DC gain. When this effect is added to the MATLAB model, the results match the SQNR predicted by SPECTRE simulations to within 2 dB.

## 6 Conclusion

The realisation of true Ambient Intelligence is achievable only if advances are made in the hardware implementation of self-contained wireless sensing. The modelling and simulation of thin-film batteries, MEMS components, RF transceivers, network algorithms and low energy analog circuitry have been described. It is shown that accurate modelling and simulation of complex multiple-domain systems are crucial for achieving reliable and efficient operation. Currently, *ad hoc* behavioural modelling and simulation augments fine-grained, computationally intensive simulations in each of the disciplines are discussed. The interface between technologies is typically modelled by a greatly abstracted representation of the neighbouring system. As technology advances to bring communication, energy harvesting, sensing and computation on the same chip, the need for modelling and simulating complex interactions between domains becomes even more acute.

## 7 Acknowledgments

The authors would like to acknowledge the Mixed Signal Circuits and Systems division of Philips Research Laboratories, Eindhoven for funding the work on High-Speed Cascaded modulators, as well as providing the measured chip. Special thanks to Dr Lucien and J. Breems for technical guidance and discussions.

## 8 References

- Weber, W., Rabaey, J.M., and Aarts, A.: 'Ambient intelligence' (Springer, New York, NY, 2005)
- Rabaey, J., Ammer, J., Karalar, T., Li, S., Otis, B., Sheets, M., and Tuan, T.: 'PicoRadios for wireless sensor networks: the next challenge in ultra-low-power design'. Proc. The Int. Solid-State Circuits Conf. (ISSCC), San Francisco, CA, 2002
- Roundy, S., Steingart, D., Frechette, L., Wright, P., and Rabaey, J.: 'Power sources for wireless sensor networks', *Lecture Notes Comput. Sci.*, 2004, **2940**, pp. 1–17
- Shah, R.C., Wietholter, S., Wolisz, A., and Rabaey, J.: 'Modeling and analysis of opportunistic routing in low traffic scenarios'. IEEE Third Int. Symp. on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, (WIOPT), 2005
- Shah, R.C., Wietholter, S., Wolisz, A., and Rabaey, J.: 'When does opportunistic routing make sense?'. *IEEE Third Int. Conf. on Pervasive Computing and Communications Workshops (PerCom)*, 2005
- Varga, A.: 'The OMNeT++ discrete event simulation system'. European Simulation Multiconf., June 2001
- Kopke, A., Willig, A., and Karl, H.: 'Chaotic maps as parsimonious bit error models of wireless channel'. IEEE Twenty-Second Annual Joint Conf. IEEE Computer and Communications Societies. (IEEE InfoCom), March 2003
- Ruby, R., Bradley, P., Larson III, J., Oshmyansky, Y., and Figueredo, D.: 'Ultra-miniature high-Q filters and duplexers using FBAR technology'. IEEE ISSCC, 2001, pp. 120–121
- Li, S.-S., Lin, Y.-W., Xie, Y., Ren, Z., and Nguyen, C.T.-C.: 'Micromechanical "hollow-disk" ring resonators'. Technical Digest, IEEE MEMS Conf., 2004, Maastricht, The Netherlands, pp. 821–824
- Lakin, K.M.: 'Thin film resonators and filters'. Proc. IEEE Ultrasonics Symp., October 1999, pp. 895–906
- Quévy, E.P., Bhawe, S.A., Takeuchi, H., King, T.-J., and Howe, R.T.: 'Poly-SiGe high frequency resonators based on lithographic definition of nano-gap lateral transducers'. 2004 Solid State Sensor, Actuator and Microsystems Workshop (Hilton Head 2004), Hilton Head Island, South Carolina, June 2004, pp. 360–363
- Rantakari, P., Kaajakari, V., Mattila, T., Kiihamäki, J., Oja, A., Tittonen, I., and Seppä, H.: 'Low noise, low power micromechanical oscillator'. Proc. 13th IEEE Int. Conf. on Solid-State Sensors, Actuators and Microsystems, Seoul, Korea, June 2005, pp. 2135–2138
- Wong, A.-C., and Nguyen, C.T.-C.: 'Micromechanical mixer-filters ("Mixlers")', *J. Microelectromech. Syst.*, 2004, **13**, pp. 100–112
- Bindel, D.S., Quevy, E., Koyama, T., Govindjee, S., Demmel, J.W., and Howe, R.T.: 'Anchor loss simulation in resonators'. Proc. IEEE MEMS 2005 Conf., Miami, FL, January 2005, pp. 133–136
- Otis, B., and Rabaey, J.: 'A 300 uW 1.9 GHz CMOS oscillator utilizing micromachined resonators', *IEEE J. Solid State Circuits*, 2003, **38**, pp. 1271–1274
- Kwok, K., and Luong, H.: 'Ultra-low-voltage high-performance CMOS VCOs using transformer feedback', *IEEE J. Solid State Circuits*, 2005, **40**, pp. 652–660
- Otis, B., Chee, Y.H., and Rabaey, J.: 'A 400 uW Rx, 1.6 mW Tx super-regenerative transceiver for wireless sensor networks'. IEEE Int. Solid-State Circuits Conf. (ISSCC), San Francisco, CA, 2005
- Lai, X., and Roychowdhury, J.: 'Capturing oscillator injection locking via nonlinear phase-domain macromodels', *IEEE J. Microw. Theory Tech.*, 2004, **52**, (9), pp. 2251–2261
- Sangiovanni-Vincentelli, A.: Defining platform-based design EE times, February 2002
- Liu, E., Demir, A., Malavasi, E., and Sangiovanni-Vincentelli, A.: 'A top-down constraint-driven design methodology for analog integrated circuits' (Kluwer Academic Publisher, 1995)
- Enz, C., Krummenacher, F., and Vittoz, E.A.: 'An analytical MOS model valid in all regions of operations and dedicated to low voltage, low current applications' in 'In analog signal processing' (Kluwer, 1995)
- Castello, R., Grassi, A., and Donati, S.: 'A 500 nA sixth-order bandpass SC filter', *IEEE J. Solid-State Circuits*, 1990, **25**, (3), pp. 669–676
- Boser, B.E., and Wooley, B.: 'The design of sigma-delta analog to digital converters', *IEEE J. Solid-State Circuits*, 1988, **23**, (6), pp. 1298–1308
- Leenaerts, D., and Bolkhoven, W.: 'Piecewise linear modeling and analysis' (Kluwer Academic, 1996)