Exploring the unknown through successive generations of low power and low resource versatile agents

Martin Andraud*†, Gonenc Berkol*, Jaro De Roose†, Santosh Gannavarapu†, Haoming Xin*,
Eugenio Cantatore*, Pieter J.A. Harpe*, Marian Verhelst†, Peter G.M. Baltus*
*TU Eindhoven, Dept. of Electrical Engineering, MSM group, The Netherlands
{M.Andraud, G.Berkol, E.Cantatore, P.J.A.Harpe, P.G.M.Baltus}@tue.nl
†KU Leuven, Dept. of Electrical Engineering (ESAT), MICAS group, Belgium
{jare.deroose, santosh.gannavaparu, marian.verhelst}@kuleuven.be

***Abstract*—The Phoenix[1] project aims to develop a new approach to explore unknown environments, based on multiple measurement campaigns carried out by extremely tiny devices, called agents, that gather data through multiple sensors. These low power and low resource agents are configured specifically for each measurement campaign to achieve the exploration goal in the smallest number of iterations. Thus, the main design challenge is to build agents as much reconfigurable as possible. This paper introduces the Phoenix project in more details, and presents first developments in the agent design.**

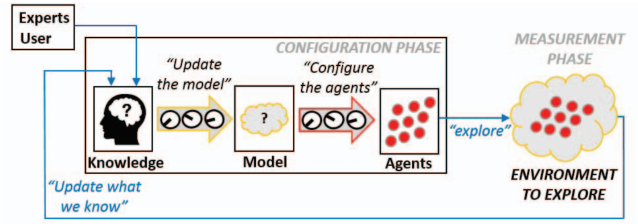

Fig. 1. general description of the Phoenix approach

I. Introduction

Even if humans have been exploring the world for centuries, many environments remain inaccessible, because they are difficult or dangerous to access. However, there is a high need for extracting information from them, which is primarily a topological map, but that can be augmented by sensory data. We would then be able e.g. to measure the pressure map inside an oil well, or to evaluate the extent of a pollution in an underground river. A recent trend is to deploy autonomous electronic sensor nodes to explore and monitor this type of environments. A large monitoring time is typically achieved through smart *ad hoc* optimization of sensing schemes. Yet, some environments are hard to reach even for such cutting edge technology, e.g. when preventing any form of communication outside the environment, and/or when the size of sensing devices must be limited to a few mm^3 (think that they must be injected via a mechanical pump for instance). In addition, it is very difficult to optimize a sensory system if the environment properties, and the signals to be sensed, are to a great deal unknown, making impossible to determine a-priori how sensory devices should behave. Hence, the exploring system struggles with a fundamental resource-information conflict. Ideally, it has to sense accurately for the whole exploration time in order to maximize the gathered information. However, the amount of energy available may not be sufficient to achieve a complete exploration with such small exploring devices. Consequently, the system must be optimized to sense more efficiently, while limiting information losses. This strategy requires to know the environment beforehand and optimize the system accordingly, which is not possible here. An alternative, as the most recent works suggest, is to reconfigure the hardware on the fly. However, in our situation, communication with the sensory devices from outside is impossible, and they do not have enough resources themselves to do such a massive data fusion. For those reasons, a new paradigm shift is required. The Phoenix project proposes a new line of technology to explore unknown and inaccessible environments. The key ideas are:

1) Multiple measurement campaigns are carried out inside the environment, with successive generations of *agent* swarms. These agents are tiny devices, able to penetrate the environment and gather information through sensors. In addition, they are very size and resource-limited, and have to operate without direct external control over software and hardware.

2) Prior to every new campaign, the generation of agents gets optimized again. This is necessary since they cannot sense all information in a single exploration due to energy constraints.

The Phoenix project innovates on the exploration approach, optimization algorithms, and on the flexible hardware needed to realize the agents, which is the focus of this work.

II. The Phoenix approach

The first target of the Phoenix project is the exploration of inaccessible environments in a fluid medium (e.g. underwater pipes, oil wells, etc.). In order to detail our proposed approach, we will start with an application example. Specifically, we would like to obtain the pressure map inside an oil well, to identify areas were the drilling can become risky due to a possible overpressure. We note that the well is mostly inaccessible, and thus a-priori information is very limited. The exploration will be performed according to the simplified description of Fig.1. Here, our goal is to build a topological map of the well and link it with pressure measurements carried out in-situ. In details, Phoenix starts with processing a user question e.g *"which is the pressure map inside the oil well?"*. A

---

[1] Phoenix has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 665347.

978-3-9815370-8-6/17/$31.00

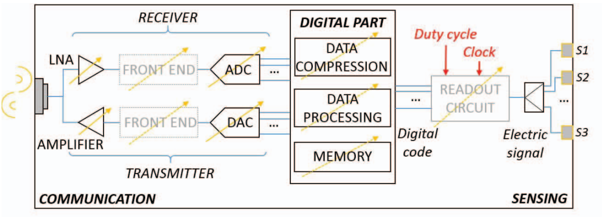

Fig. 2. Block diagram of the envisioned Phoenix agent

TABLE I
PERFORMANCE COMPARISON OF UNDERWATER MODEMS

Ref.	**Freq. (kHz)**	**Transmitter Power (W)**	**Transmission Distance (km)**	**Receiver Power (W)**	**Bit Rate (bps)**
[1]	18	2	0.05-0.5	0.025	600
[2]	25	NS	<0.5	NS	133
[3]	9-14	NS	3	NS	1200
[4]	85	0.012	0.1	0.024	1000
[5]	50	NS	<0.02	NS	80
Our target	1000	125.10^{-6}	0.01	20.10^{-6}	20000

first generation of agents has to be configured to perform a first exploration. Since the environment has not been explored yet, Phoenix proposes to configure the agents based on available knowledge, coming from the user or various experts (e.g. an initial guess about the fluid composition, the well volume, etc.). This knowledge is analysed and Phoenix creates a model of the environment, limited at first but that will be gradually improved with the results of future explorations. Then, Phoenix enters a loop with two phases. In the *measurement* phase, the agents are deployed in the environment to collect data, retrieving (part of) them after exploration. Then, the *configuration* phase, contains several steps. First, exploration data is analysed and used to update the current knowledge, further used to improve the environment model. If this model is not accurate enough, a new generation of agent is configured, optimized to obtain the missing information. Another measurement is conducted, and the overall loop is repeated until the model of the environment is sufficiently accurate to answer the user's question. Consequently, a critical constraint for Phoenix is to have an agent as configurable as possible. We refer to this as *versatility*. The more versatile the agent is, the more possibilities Phoenix has to optimize it for each measurement phase, and hence the less iterations are needed to obtain the most accurate answer. The agent will be detailed in the next sections, together with first developments concerning communication, sensing and data processing techniques.

III. THE PHOENIX VERSATILE AGENT

From a hardware perspective, the agent has to be optimized even for tasks that are unknown at design time. In addition, the agent will be extremely miniaturized, and so very limited in energy and resources. Our guideline is then to perform each operation in the most energy-efficient way. Following up on the example presented in section II, but without loss of generality in the approach, other design challenges include:

- **Communication**: data exchange between agents is a key to localize them and obtain the environment map. But communicating in a fluid medium imposes specific requirements.
- **Sensing**: since the agent is configured specifically for each campaign, the exact number of sensors and sensed parameters are unknown at design time. Thus, embedding a dedicated front-end for each possible sensor is not feasible.
- **Digital part**: an enormous amount of data may be collected by the agent throughout exploration, although most relevant information may be sparse and available memory very low.

Tackling all these challenges, our envisioned agent block diagram is depicted in Fig.2. As it can be seen, the agent is composed of three main parts, to enable communication, sensing and digital processing. The associated solutions that we propose will be detailed in the next subsections.

A. Communication

The communication between agents will be a key feature to enable their localization and build a topological map of the environment. However, communicate inside a fluid imposes specific constraints, especially in our context. Indeed, the communication distance is very small ($10m$), and power consumption must be reduced at maximum to achieve the best energy efficiency. A first step is to select the most appropriate way to communicate in a fluid medium.

a) State of the art of underwater modems: Since current state-of-the-art targets mainly underwater communication systems, it will be the focus of this section. However, more complex scenarios including liquids of different viscosity and density are also in the project scope. First of all, underwater data transmission has been investigated with radio-frequency (RF), optical, and ultrasonic waves [6]. In the context of Phoenix, ultrasonic waves have been selected because they have the smallest path loss, and the speed of sound is relatively small which enables to design low speed front-end systems. Furthermore, thanks to the advancements of piezo micro-machined transducers, driving voltages could be made compatible to CMOS technology and transducers can be miniaturized [7]. To the best of our knowledge, an integrated circuit solution that focuses solely on underwater communication has not been published yet. Therefore, we focus here on research modems implemented with discrete components, listed in Table I (see [8], [9] for a detailed discussion). Commercial products have also been developed to achieve large communication distances, but they consume a large amount of power. Concerning research modems, the state of the art in terms of power consumption is presented in [4], where $12mW$ transmitter and $24mW$ receiver powers are used to achieve $0.1km$ communication distance with $1000bps$.

b) Derivation of first specifications for Phoenix: In Phoenix, the ultrasound transducer will be specifically developed for the project, due to the small-size required for the agent, thereby increasing the communication frequency up to $1MHz$. Our objective is to achieve at least 10^{-3} bit error rate for On-Off Keying (OOK) modulated signal with those constraints. The attenuation of ultrasound waves in underwater occurs due to spreading loss and frequency dependant absorption loss. Also, thermal noise is considered to

TABLE II
PERFORMANCE COMPARISON BETWEEN VERSATILE SENSOR INTERFACES

Ref.	Sensor type	Sensor range	Output type	Resolution (bits)	Bandwidth	Power
[12]	Capacitive Resistive	C:10-100pF R:not reported, tested at $22k\Omega$	Analog	C:16.6 R:11	10Hz	$53\mu W$ (typical)
[13]	Capacitive	2.5-75.3pF	Digital	13.3 (worst case)	125Hz	$160nW$
[14]	Capacitive	0.54-1.06pF	Digital	12.6	625Hz	$10.3\mu W$
[15]	Capacitive	1-15pF	Digital	8 to 10	1-62.5Hz	$7\mu W$ (8b,5Hz) $10\mu W$ (9b,5Hz)
Our target	Capacitive Resistive Conductivity	C:1pF-20pF R:$20k\Omega$-$10M\Omega$ Cond:1-100mS/cm	Digital	6 to 10	1Hz to 100Hz	$< 1\mu W$

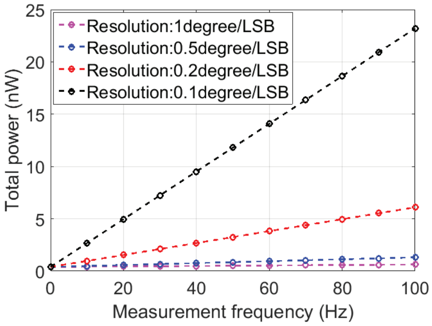

Fig. 3. Power vs. measurement frequency plot for the resistive temperature sensor interface

be dominant underwater for frequencies greater than $100kHz$ [10]. Considering a case study with $1MHz$ frequency, the voltage efficiency at 10m distance can be estimated around $-110dB$, which corresponds to $5.6nV/\sqrt{Hz}$ required noise density at $200kHz$ bandwidth at the receiver input, assuming a SNR of $16dB$, which is sufficient to achieve a target BER with OOK. The required noise density could be achieved by using $17\mu A$ biasing current (assuming a simple amplifier in the receiver), leading to a $20\mu W$ power consumption with 1.2V supply. At the transmitter side, we consider that the most of the power is used to charge the transducer capacitor, assumed to be around 500pF for piezo-electric transducers [11]. We expect a power consumption around $0.125mW$ with a driving voltage of $5V$ (which is compatible with most commercial low-voltage CMOS technologies) if the transmitter is active for $100\mu sec$ per second. Our target is summarized in Table I.

c) Comparative of existing and target specifications: As we can see, none of the existing work is fitting to our goal, either in terms of power consumption and transmission distance, or in terms of frequency. As a result, our main objective is to design integrated circuits, in order to develop ultra low-power sensor nodes suitable to our application.

B. *Versatile sensor interface*

Usually, the sensing system is tailored for one particular sensor with a given accuracy and speed to maximize system performance. However, in our context, the type and number of sensors are unknown at design time. The idea developed for Phoenix is then to sense many parameters with a unique interface, as illustrated in Fig.2. All sensors share the same readout circuit, that stimulates each sensor and digitize the measurement. In terms of versatility, the interface must comply with sensors of different types, all with their own ranges and sensitivities. Referring to our example, we could use temperature, pressure and conductivity (i.e. salinity) sensors, to extract the fluid properties and sense the pressure in-situ, as well as an accelerometer to obtain a more accurate localization. Sensing could then be performed by a resistor for temperature, by capacitors for pressure and acceleration, and by an impedance sensor for salinity. In a more generic perspective, those three types must be taken into account in our final implementation, since they gather a vast majority of existing sensors. Also, the interface must be configured for each measurement campaign and application scenario. In our example, we could focus on an accurate pressure measurement, requiring high performance, or on determining the topology which would require a low-power sensor mode. Moreover, energy efficiency is of primary concern, with a total target power consumption below $1\mu W$.

a) State of the art: Due to design complexities, power inefficiency or large silicon area, there is no existing sensor interface that can simultaneously cover resistive, capacitive and conductivity measurements. Existing designs that show part of this versatility are listed in Table II. It can be concluded that they either lack of versatility to tune their performance [13]–[15] or have the required versatility but at the expense of too high power consumption [15] for Phoenix. Therefore, new concepts and hardware architectures will be developed.

b) First case study: A first versatile interface is currently being designed in a 65nm CMOS, focusing on differential resistive temperature measurements. The versatility is twofold. First, the measurement speed can be tuned from 1Hz to 100kHz with an external clock. Besides, we use the oversampling and averaging method to increase the measurement accuracy while scaling the power consumption. In order to keep an optimal power efficiency, the duty-cycle of the measurements can also be tuned, to perform less measurements if each of them requires a larger amount of power. The performance of the interface, focusing on low-speed applications, is summarized in Fig.3, where data is extracted from post layout simulations. As it can be seen, the power consumption scales linearly with the measurement frequency, and the temperature resolution can be tuned from 1°C$/LSB$ (Least Significant Bit of the ADC) to 0.1°C$/LSB$ by tuning the oversampling ratio.

C. *Data compression*

During exploration, the agent will have to store recorded data streams into a non-volatile flash memory. The possible amount of collected data is enormous. Following up on our example, we have for instance data from communication and all sensors. Because of extremely low-power constraints, we have to minimize the computational energy used to extract and process data, and the energy spent on storage itself. Our strategy is thus to compress data inside the agent so that the effectively stored data is minimized. Then, the energy consumption is traded off against the amount of information loss tolerated in the compression, depending on the scenario envisaged (type of sensor, accuracy needed, etc.).

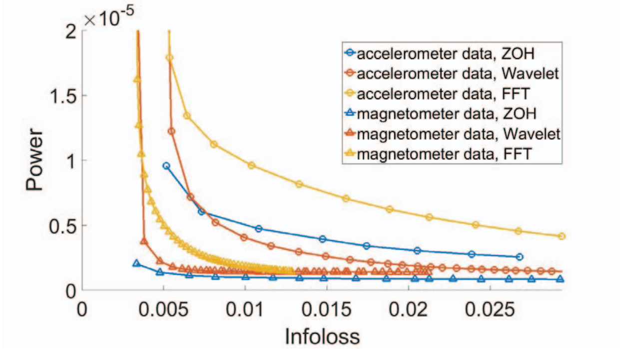

Fig. 4. Data compression from accelerometer and magnetometer data

a) Low footprint data extraction and compression techniques: The available tuning parameters are: the sensory data sampling frequency f_s, the ADC resolution N_{ADC} and the compression algorithm threshold Δ. Moreover, several lossy compression algorithms can be used and need to be compared:
Zero order hold: we compare the current sample to the previous, and only store it if the two data points differ more than a predefined threshold Δ_{ZOH}.
Wavelet transform compression: we first perform a wavelet transform on a sample window. The transformed values are subsequently compressed: only wavelet transform coefficients above a predefined threshold Δ_{wav} are stored.
Fourier transform compression: Similarly to the Wavelet algorithm, a Fast Fourier Transform (FFT) is performed, after which only FFT coefficients above Δ_{FFT} are stored.

b) Information loss and total energy calculation: Information loss for different settings is compared with the Percentage Root-mean-square Difference (PRD) defined as:

$$PRD = \frac{\sqrt{\sum_{n=1}^{k}(X_n - Y_n)^2}}{\sqrt{\sum_{n=1}^{k} X_n^2}} = \frac{RMS(ErrorSignal)}{RMS(Signal)} \qquad (1)$$

where X is the raw sensory data and Y the compressed signal. The energy usage is assessed by building a complete energy model of the data acquisition, compression and storage hardware regarding f_s, N_{ADC}, C and Δ. Table III shows the energy estimations, assuming an implementation in a 90nm CMOS chip. The energy per sample needed for the different compression algorithms is also derived, based on the number of additions and multiplications required per sample. It is then clear that in most cases, the storage energy will be dominant, indicating the need for a high compression rate.

c) Simulation results: The trade-off is analysed on Phoenix data collected with the INCAS3 motes [16]. Fig. 4 shows the PRD-Power plot of accelerometer and magnetometer sensor data. The f_s and N_{ADC} are held constant at 309Hz and 10 bit/sample, while Δ is swept. We can observe that in the accelerometer case, the wavelet transformation compression gives better performance, which can be contributed to spikiness of accelerometer data. In the magnetometer case, Zero order hold performs better due to the low data resolution.

TABLE III
ENERGY REQUIRED FOR EACH OPERATION

Operation	Energy Consumption
Analog Front end (E_{AF})	1nJ/sample
Analog/digital conversion (E_{ADC})	1pJ/sample
Addition (16-bit)	2pJ
Multiplication(16-bit)	10pJ
Zero Order Hold (E_{DC})	2pJ/sample
Wavelet transform (N=256) (E_{DC})	200pJ/sample
Fourier transform (N=256) (E_{DC})	48pJ/sample
Memory write (Flash) (E_S)	2.5nJ/bit

IV. CONCLUSION

This work focused on the first hardware developments carried out in the Phoenix project, that requires extremely versatile exploring agents. After explaining the challenges and envisaged solutions, we have illustrated our current work and results on communication, sensing and data processing parts.

REFERENCES

[1] J. Wills, W. Ye, and J. Heidemann, "Low-power acoustic modem for dense underwater sensor networks," in *Proceedings of the 1st ACM international workshop on Underwater networks*, 2006, pp. 79–85.
[2] R. Iitis, H. Lee, R. Kastner, D. Doonan, T. Fu, R. Moore, and M. Chin, "An underwater acoustic telemetry modem for eco-sensing," in *Proceedings of OCEANS 2005 MTS/IEEE*, 2005, pp. 1844–1850.
[3] E. M. Sözer and M. Stojanovic, "Reconfigurable acoustic modem for underwater sensor networks," in *Proceedings of the 1st ACM international workshop on Underwater networks*, 2006, pp. 101–104.
[4] A. Sanchez, S. Blanc, P. Yuste, and J. Serrano, "A low cost and high efficient acoustic modem for underwater sensor networks," in *OCEANS 2011 IEEE-Spain*, 2011, pp. 1–10.
[5] B. Benson, G. Chang, D. Manov, B. Graham, and R. Kastner, "Design of a low-cost acoustic modem for moored oceanographic applications," in *Proceedings of the 1st ACM international workshop on Underwater networks*, 2006, pp. 71–78.
[6] I. F. Akyildiz, D. Pompili, and T. Melodia, "Underwater acoustic sensor networks: research challenges," *Ad hoc networks*, vol. 3, no. 3, pp. 257–279, 2005.
[7] Z. Deng, T. J. Carlson, H. Li, J. Xiao, M. J. Myjak, J. Lu, J. J. Martinez, C. M. Woodley, M. A. Weiland, and M. B. Eppard, "An injectable acoustic transmitter for juvenile salmon," *Scientific reports*, vol. 5, 2015.
[8] S. Sendra, J. Lloret, J. M. Jimenez, and L. Parra, "Underwater acoustic modems," *IEEE Sensors Journal*, vol. 16, no. 11, pp. 4063–4071, 2015.
[9] B. Benson, Y. Li, R. Kastner, B. Faunce, K. Domond, D. Kimball, and C. Schurgers, *Design of a low-cost, underwater acoustic modem for short-range sensor networks*. IEEE, 2010.
[10] M. Stojanovic and J. Preisig, "Underwater acoustic communication channels: Propagation models and statistical characterization," *IEEE Communications Magazine*, vol. 47, no. 1, pp. 84–89, 2009.
[11] J. Johansson and J. Delsing, "Microelectronics mounted on a piezoelectric transducer: Method, simulations, and measurements," *Ultrasonics*, vol. 44, no. 1, pp. 1–11, 2006.
[12] J. Zhang, J. Zhou, and A. Mason, "Highly adaptive transducer interface circuit for multiparameter microsystems," *IEEE Transactions on Circuits and Systems I, Regular Papers*, vol. 54, no. 1, pp. 167–178, 2007.
[13] H. Ha, D. Sylvester, D. Blaaw, and J.-Y. Sim, "A 160nw 63.9fj/conversion-step capacitance-to-digital converter for ultra-low-power wireless sensor nodes," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*. IEEE, 2014, pp. 220–221.
[14] H. Danneels, K. Coddens, and G. Gielen, "A fully-digital, 0.3v, 270nw capacitive sensor interface without external references," in *Proc. of European Solid-State Circuits Conference ESSCIRC*, 2011, pp. 287–290.
[15] W. Bracke, P. Merken, R. Puers, and C. Van Hof, "Ultra-low-power interface chip for autonomous capacitive sensor systems," *IEEE Transactions on Circuits and Systems I, Regular Papers*, vol. 54, no. 1, pp. 130–140, 2007.
[16] E. Talnishnikh et al., "Motes for environment mapping, NL patent 2011892."