Architecting the Human Intranet

Jan M. Rabaey
EECS Department
University of California
Berkeley, USA
janrabaey@berkeley.edu

Ana Claudia Arias
EECS Department
University of California
Berkeley, USA
acarias@berkeley.edu

Rikky Muller
EECS Department
University of California
Berkeley, USA
rikky@berkeley.edu

***Abstract*—**Equipping us humans with the necessary tools to interact with, survive, and prosper in a rapidly changing world may require us to intimately adopt some of the same technologies that are causing some of these changes. Various wearable devices have been or are being developed to do just that. To be effective, functionality cannot be centralized and needs to be distributed to capture the right information at the right place. This requires a human intranet, a platform that allows multiple distributed input/output and information processing functions to coalesce and form a single application. How to effectively do so in light of the many challenges from an efficiency, usability, and effectiveness perspective is the focus of this paper.

Keywords—wearables, implants, body-area networks, flexible electronics, ultra-low energy

I. INTRODUCTION

The world as we know it is going through major upheavals: climate change, pandemics and technology-induced societal changes are upsetting our world-picture with no real end in sight. Hence, an extremely relevant question is how 'we humans' are going to cope with all these rapid evolutions. One plausible answer is for us to use those same technologies that we are unleashing on the world to evolve ourselves, and to equip us with the necessary tools to interact with, survive, and prosper in spite of (or in light of) these changes.

Various wearable and implantable devices have been or are being developed to do just that. However, their potential to create a whole new set of human experiences is still largely unexplored. To be effective, functionality cannot be centralized and needs to be distributed to capture the right information at the right place. This requires a *human intranet* (*HI*) a platform that allows multiple distributed input/output and information processing functions to coalesce and form a single application [2]. In this paper, we focus on the integrative aspects of such an intranet, a function that is complicated by the extreme energy and form-factor constraints imposed on the composing wearable/implantable devices. Analyzing various application from that perspective, it becomes clear that the human intranet should be capable of dealing with changes in both the environment and the body itself, and learn from experience. Moreover, it should be able to do so on a continuous base. To inspire how this may be accomplished, it is worthwhile studying the biological proof-of-concept, that is the human nervous system itself. In a way similar to what we intend for the intranet, it collects information from a range of heterogeneous senses, analyzes that information through multiple stages of processing, and ultimately plans and directs action (motor function). The multiple feedback loops governing that behavior are illustrated in Figure 1 (from [3]). Most essential is a separation between inner and outer loops. The former ensure that the body itself performs correctly by observing metabolic and physiological parameters, and by taking autonomous actions to correct or adjust. The outer loops on the other hand serve to observe and interact with the environment around us. Information (represented as patterns) is gradually fused into ever-larger patterns representing higher layers of abstraction. Similarly, after evaluating various output options, the selected high-level pattern is gradually decomposed into localized actions. One of the most compelling features of this architecture is its self-healing nature and the capability to adapt and adjust based on internal and external circumstances. These are exactly the functions *HI* is attempting to assist, complement and augment.

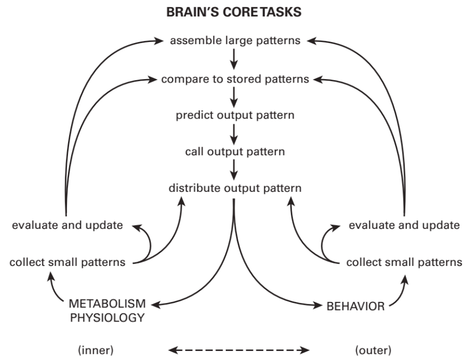

Figure 1. The neural system and its feedback loops (from [3]).

In the coming sections, we briefly address the various components that make up such a system, starting with the network itself, the information sensing and acquisition, the pattern classification and processing functions, and finally the actuation. To help make the discussion concrete and to derive clear specifications, we present some examples of human-intranet applications/configurations first.

II. COMPELLING HI APPICATIONS

Human Intranet applications generally fall into two categories. Many systems fall into the wellness and health space, such as closed-loop control of blood sugar, or the correlation of blood oxygenation with breathing rhythms. Other applications fall into the domain of augmentation such as enhanced vision and exoskeletons. The exemplary applications discussed in this section are only a fraction of what is possible.

A. *Hearables and Hearing Aids*

Hearing Aids have been an essential part in addressing hearing loss. Yet, these medical devises are expensive, require careful tuning, and often are unsatisfactory. With the advent of all styles of earbuds, the gap between medical and consumer devices, intended to provide enhanced audio enjoyment, is rapidly fading [4]. Advanced signal processing embedded in these devices helps to customize the response based on the estimated audio propagation channel between source and the brain, and to cancel various sources of interference. Wirelessly connected to a mobile phone or smart watch, sensing (microphones) and processing can be distributed to capture the best possible 3D audio capture or to perform computation where energy is available. The system can be informed of intent or context by auxiliary sensors such as eye tracking and motion. Going one step further, in-ear EEG sensors can gauge parameters such as stress, effort or emotion, to further optimize the hearing response or simply serving as another biometric parameter to be used in wellness observation (Figure 2).

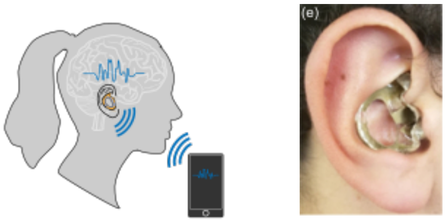

(a) (b)

Figure 2. In-ear EEG sensor (earbud) provides brain-computer interface (adopted from [17]).

As observed in [4], all these elements together are leading to new classes of hearables customized for a broad range of functions (voice, music, concert, sports, enhancement, and correction). Making this happen will require various aspects of the Human Intranet to come true: ultra-low energy heterogeneous sensing, robust networking and distributed intelligence. The network needs to effectively support audio data streams (> 750 Kbit/sec) and very-low latency (~ 1 msec) response times in overlay with low and medium rate contextual information channels. While most processing can be performed at the hub Node (the mobile), localized processing in the sensors and the earbud may be necessary to address latency, bandwidth and robustness concerns. The envisioned intelligence (see e.g. [5]) may easily supersede what can be performed in the devices of today.

B. *Smart human-in-the-loop prostheses*

The technology of prostheses has been steadily but slowly improving over the past years. Yet, most are still heavy, energy-hungry and, most importantly, cumbersome to use. Essential to satisfactory behavior is a close interplay between the human and the device. While ideally, the intent of the person would be derived directly from the brain, EMG signals captured from the upper arm or leg serve as a surrogate [6]. Unfortunately, these signals are noisy, and time-variant depending upon the context, conditions and wear. Also, closing the feedback on how successful the prosthesis is in performing the task is mostly visual. Performing a grasping task requires users to employ 90% of their visual attention! Considering other sensor modalities such as motion, pressure and stress level can substantially improve the experience, but is highly individual and may vary over time and circumstances. Distributed learning can help to address these concerns but should be lightweight and individualized. However, this is just the start. A wider range of inputs could create a tighter connection between intent and outcome, or, between human and prosthesis (Figure 3). Similar considerations hold for applications such as exoskeletons. Compared to the hearing applications, data rates are substantially lower here, but robustness is substantially more critical. Privacy and security also play a larger role.

C. *Brain-Machine Interfaces (BMI)*

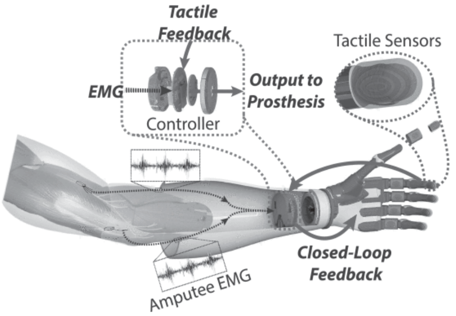

Figure 3. Prosthetic control shared between user intention and robotic automation (adopted from [1]).

By far the most ambitious of all envisioned HI scenarios, BMIs offer a direct interface to the human computer (the brain) enabling both read (output) and write (input) functionality. This field has made gigantic steps forward over the past decade, primarily with regards to the biological-electrical interface. Academic and industrial endeavors are creating solutions that offer thousands of channels with varying degrees of invasiveness ([18,19]). Yet, we are far from offering solutions that could be used in an everyday setting, or that effectively close the loop between brain read and write(actuation/stimulation). This explains why BMI is a most alluring HI application. Real deployment will require distributed operation, sensor fusion, efficiency, conformity, safety, privacy, security, and definitely smartness. And perhaps most of all – ethical guidelines.

III. THE NETWORK

While all parts of HI play equally crucial roles, the network is truly at the core of the system, similar to how the nervous system integrates all parts of the human body. Enabling robust, reliable and efficient communication between a wide variety of devices requiring a very wide range of data rates is a major

---

Major funding for thus work has been provided by the SRC/DARPA Sonics and CONIX centers; the DARPA HYDDENN and VIP programs, and the member companies of the UC Berkeley BWRC center.

challenge. The IEEE 802.15.6 WBAN (Wireless Body Area Network) standard was established to that effect. But it has some intrinsic limitations in addressing the heterogeneity of the network nodes in terms of activity, data rates (from less than bit/s to multiple Mbit/s), energy availability, form factor, and location on, around or inside the body. To address these concerns, an alternative model optimized for communication around a living body is proposed (it can be considered as an extension of IEEE 802.15.6). At its core are the following propositions:

- To avoid the unpredictive nature of RF communication around the body, the body itself is used as the communication channel. This allows for efficient non-LOS (Line of Sight). The concept of human body coupled (HBC) communication is actually quite old [7], it is already a part of the IEEE 802.15.6 standard, and has gathered some new attention recently [8]. However, relying on the traditional low-frequency capacitive communication approach limits the bandwidth and hence the maximum data rates achievable. When operating at higher carrier frequencies (~ 500 MHz), information propagates as surface waves [9] enabling a substantially higher bandwidth (> 100 MHz) and data rates up to 30 Mbit/sec [10].

- Efficient and predictable communication is more easily achieved in a synchronous scheme, in which access to the media is governed on a temporal base (such as TDMA). This however requires synchronization between the nodes, which can be very expensive especially for the small sensor nodes. Fortunately, the body already features a global synchronization signal that can be acquired at very low energy cost (nWs): the heartbeat. A heartbeat-based Media Access Control (MAC) protocol, which is fully decentralized, has been proven to provide reliable and predictable communication over many operational conditions [11].

- Providing robustness and reliability requires redundancy. In addition, power limitations on the smallest nodes limit their physical distance reach. Both arguments point to a star-mesh network, as the preferred network topology (Figure 4). It defines two types of nodes – hubs and leaves. The former offer greater computing capabilities and energy availability, whereas the latter are a lot more constrained and energy-frugal. Network topology is dynamic, and can change with the addition/removal of nodes and application needs.

- Most of the application traffic of the HI is periodic and predictable. Similarly, variations in the communication channels are often the result of behavioral conditions (such as sitting, walking, running, reaching, hugging, etc.). As these are recurring, they can be learned and can be used to optimize the network settings dynamically. Information about these operational conditions can be obtained from channel measurements, but, even better, they can be inferred from information extracted by sensors in the network itself observing kinematics and biometric parameters [12].

Truth be told, this whole concept is still in its infancy. Networks around the human body are still very anchored in established wireless technologies and protocols, leading to inefficiency, frailty, and brittleness. Even more, they are vulnerable to external interference and intrusion. Getting inspiration from the way biological networks operate and adapt can help us to envision systems that seamlessly operate, adapt and survive for a lifespan.

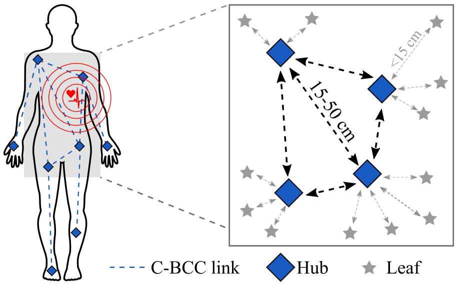

Figure 4. Human Intranet Network Architecture using heartbeat synchronization.

IV. INFORMATION ACQUISITION AND ACTUATION

Technology scaling undoubtedly is one of the major enablers of the HI concept. Over less than two decades wireless sense-and-compute nodes have scaled by more than 3 orders of magnitude in size and energy efficiency. While smaller transistor sizes (Moore's law) have surely helped, most of the scaling in the node size is due to "More than Moore" (MtM) scaling: advanced packaging and integration techniques that allow diverse technologies, such as sensors or energy supplies, to be assembled together with more traditional silicon devices for computation and memory in a 3D form-factor [13]. As a result, fully integrated "smart dust motes" of 1 mm3 in size, as originally envisioned in 1997, have now been demonstrated [14]. Yet, to be truly unobtrusive, implanted nodes must be even smaller, and extra orders of magnitude of scaling should be pursued. In an ideal scenario, physical motes may become equivalent in size to a biological cell, improving the information flow and reducing the rejection by the body. This requires innovation and creativity. The "neural dust" [15] concept is an example of such.

For non-implanted nodes worn on the body, conformity and unobtrusiveness are an essential factor. We have witnessed a rapid progress over the past few years in the development of flexible devices for sensing, communication, energy harvesting and energy storage, using techniques such as thin-film and printing technologies as well as various forms of smart fibers and patches.

There are ample examples showing how this fundamentally changes the capabilities and span of wearable technologies. Take, for example, the case of the smart prosthesis. Current myoelectric prostheses use only a few electrodes to capture the EMG signals and suffer from long-term consistency. Various approaches to address this have been considered, as show in Figure 5. They range from tight arm-straps to address signal consistency (a) to printed flexible arrays of electrodes offering conformity and redundancy (b). In the end though, the only long-term solution may be having arrays of electrodes implanted below the skin, delivering long-term consistency (c).

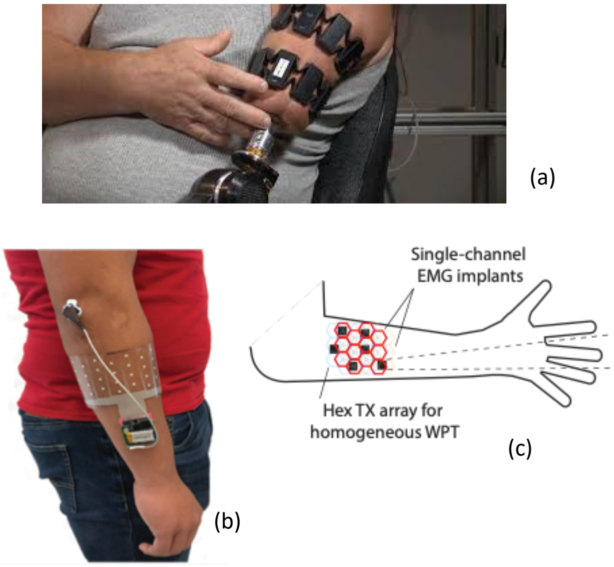

Figure 5. Various forms of EMG electrode arrays. (a) surface electrodes (myo-band); (b) flexible printed surface electrode arrays; (c) subcutaneous electrode arrays with wireless powering.

One of the prime challenges in the design of these miniature interfaces is the MtM integration – that is, a robust, reliable, long term interface between the sensing or actuating components – which are very often organic – and the data acquisition, most often realized as integrated circuit(s). The only real way to come to effective solutions is to co-design the whole system from the start, envisioning how components will be placed, how they will be interconnected and how the whole system will be encapsulated. The amount of innovation and creativity witnessed in this domain is truly mind-boggling. The design of the earpiece already shown in Figure 2 is a perfect example of this, as shown in Figure 6, which illustrates how the device was originally custom molded (c) to an individual user. Researchers then used a database of human ear measurements and recorded data from custom molded devices to inform the design of the user-generic device shown in (d).

V. DISTRIBUTED INTELLIGENCE

As stated earlier and shown in Figure 1, closed-loop behavior is of essence in the operation of the HI. The Human Intranet operates in a dynamic world subject to both slow evolution and extremely fast changes - both in the surrounding environment and in the Intranet itself - in activity, conditions, composition, and resource needs and availability. Therefore, the Human Intranet should be constructed as an adaptive and evolutionary system that combines local decision making with centralized global learning and optimization performed in hub nodes or possibly even in the cloud. This approach, in which intelligence is both global and distributed, is essential to deal with issues of latency and single points of failure, while avoiding the trap of many distributed entities with limited knowledge trying to resolve a global function. There are strong incentives to move the analysis of sensor data and the generation of actuation patterns for motor functions as close as possible to the end points. Again, much inspiration can be gained from the biological equivalent. One of leading principles that has guided the evolution of the brain can be quoted as follows: "Send only what is needed and send it at the lowest possible rate [3]." This design principle has, for instance, guided the architecture of the visual and olfactory systems in animals and humans, in which the extraction of the small patterns (features) is performed inside or right next to the sensory arrays ("in-sensor computing"). The same holds for some motor control. For instance, spinal circuitry plays an important role in the control of movement. The same principles hold for HI. Transmitting a bit may be equivalent to many thousands of gate operations. Hence, it makes perfect sense to reduce the communication rate by introducing some local processing, while being aware that the computations that can be performed are extremely constrained by the available energy.

Looking at the nature of the computation required, tasks such as signal conditioning, transformation and feature extraction can easily be performed by dedicated circuitry close to the

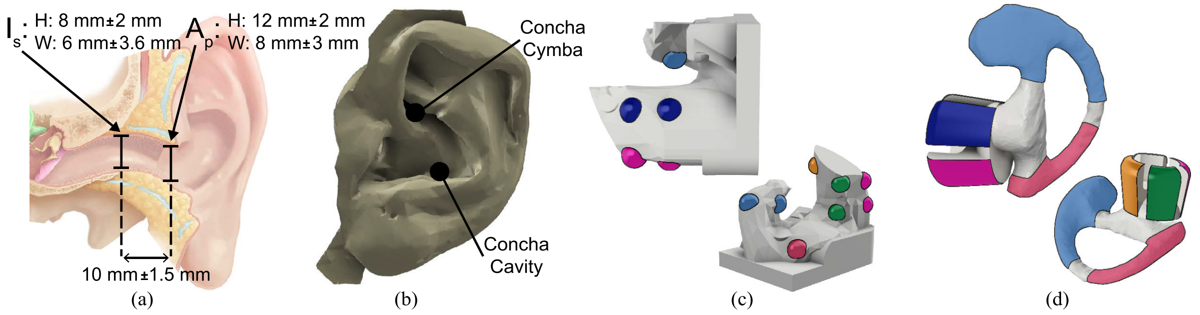

Figure 6. (a) Ear canal with average measurements; (b) High-resolution subject ear scan; (c) Earpiece custom-fit for (b) shown from two angles. Correlated electrode pairs are highlighted by the same color;. (d) User-generic earpiece. Electrodes highlighted with corresponding group colors from (c). (Adopted from [17]).

sensor/actuators. However, understanding, reasoning and decisioning require more advanced functionality, a major fraction of which would be learning-based (large patterns). Over the past few years, we have witnessed some major advances in the realization of learning-based computing for embedded applications. Yet, the state-of-the-art is still miles away from human-like intelligence, which is based on continuous learning and learning by equivalence, and often requires very little data to get started. As every human is different, no one-fit-all solution works. Customization of the processing often leads to simpler and more effective solution. In short, artificial intelligence when applied to HI has to be plastic, capable of continuous evolution and be as simple as possible. This requires some fundamental rethinking of how computing is performed, and neither GPUs or massive DNNs are the answer.

The answer may again be found in nature. Chemical computing using proteins and strands of DNA, which is at the core of how neurons communicate may be one example, delivering superb energy efficiency at the expense of speed. Another option is to compute with patterns rather than numbers, similar to the brain (as insinuated in Figure 1). It builds on the observation that long randomly-chosen vectors are (almost) orthogonal. Any deviation from orthogonality means that there is a relation between the patterns encoded in the vectors. To match, it is sufficient that two vectors are "similar" (measured using some norm, such as the Hamming Distance). High-dimensional computing (HDC) first maps incoming sensory data into high-dimensional spaces, and then proceeds to encode temporal and spatial information by algebraic operations on these vectors. The resulting pattern can be stored in memory, or can be compared against patterns that were already observed. This can be repeated, creating patterns at different levels of abstraction [16].

This computational approach comes with some interesting properties, which make it very attractive for ultra-low-power operation while supporting plasticity: (i) it is statistical and robust against errors and variations; hence, it can operate at low signal-to-noise ratios. (ii) as it is based on algebraic operations on vectors, it provides transparency and allows reasoning about the results; (iii) it learns quickly – often a few training sessions are sufficient; (iv) it can learn incrementally and continuously; (v) it computes in superposition – that means that data is superimposed on top of each other, and can be queried with a single operation; and (vi) the central element is the associative memory, which performs search and match. Association is an important part of the mammal brain as well.

An example of this computational approach, as applied to the smart prosthesis, is shown in Figure 7. EMG signals (obtained using the electrode array shown in Figure 5) and fused with inertial data provided by accelerometers, is projected into a high-dimensional space, transformed into a single vector capturing the spatio-temporal information using simple algebraic operations, and stored (learned) into an associative memory. Classification of gestures now just requires finding the closest prototype [6].

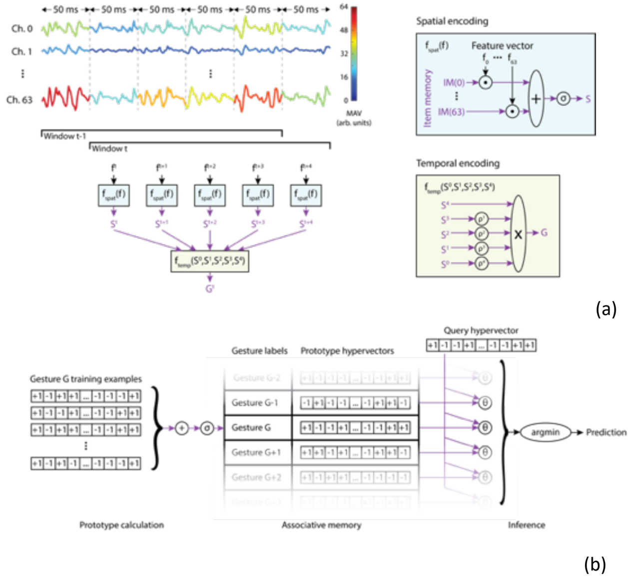

Figure 7. Gesture recognition using high-dimensional computing; (a) encoding; (b) inference.

This is just one example of how nature-inspired paradigms can help provide computational solutions to HI tasks. Many others are being pursued today, such as spiking neural nets, reservoir computing and Boltzmann Machines. The exciting part is that these approaches can benefit in a huge way from the novel device, memory and integration technologies emerging today under the MtM paradigm.

VI. PERSPECTIVES

This paper (and presentation) presents a vision on how physical and biological systems are on a convergence path. The realization of this vision to its full extent will take many more decades. It requires innovation on many fronts including interfaces, processing, communications and energy provisioning. Today, though, we already see instances emerge that give a glimpse of what may be possible in the long term: with the advent of the smart watch, wearables are becoming common place; wearable medical devices such as smart patches are making their inroads; hearing aids are starting to integrate multiple sensors and incorporate artificial intelligence; AR/VR glasses help to transform the way humans see the world around them; and brain-machine interfaces have been demonstrated in lab and clinical settings. At the same time, these systems are mostly standalone and ad-hoc, and do not profit from the wealth of information that could be available in a shared platform. It is only when an open platform such as the Human Intranet is adopted that we will see its full power, serving as a partner to our human biological systems.

ACKNOWLEDGMENT

The authors gratefully acknowledge the many contributions of students, postdocs and colleagues in making the Human Intranet come alive. More specifically: A. Araujo, R. Bennarouch, F. Burghardt, A. Cathelin, O. Facklam, A. Frappe, A. Moin, A. Kaiser, R. Kaveh, J. Doong, Y. Khan, A. Menon, A. Rahimi, F. Solt, J. Ting, A. Thielens, G. Tochou, N. Yamamoto, A. Zhou, and many others. Thanks to the sponsors of the Berkeley Wireless Research Center.

REFERENCES

[1] L. Osborn, R. Kaliki, A. Soares and N. Thakor, "Neuromimetic event-based detection for closed-loop tactile feedback control of upper limb prostheses," *IEEE Transactions on Haptics,* vol. 9, no. 2, pp. 196-206, 2016.

[2] J. Rabaey, "The Human Intranet - Where Swarms and Humans Meet," *IEEE Journal on Pervasive Computing,* vol. 14, no. 1, 2015.

[3] P. Sterling and S. Laughlin, Principles of Neural Design, MIT Press, 2017.

[4] D. Carrington, "2020 And the Future of Hearables," Forbes, May 2020. [Online]. Available: https://www.forbes.com/sites/forbestechcouncil/2020/05/21/2020-and-the-future-of-hearables/?sh=1e15b1647b84.

[5] D. Wang, "Deep Learning Reinvents the Hearing Aid," *IEEE Spectrum,* December 2016.

[6] A. Moin, et al., "A wearable biosensing system with in-sensor adaptive machine learning for hand gesture recognition," *Nature Electronics,* vol. 4, no. 1, pp. 54-63, 2021.

[7] T. Zimmerman, Personal Area NetwqorksL Near-Field Intra-Body Communication, Cambridge: MS Thesis, MIT, 1995.

[8] S. Sen, S. Maity and D. Das, "The Body is the Network," *IEEE Spectrrum,* vol. 57, no. 12, 2020.

[9] R. Benarrouch, A. Thielens, A. Cathelin, A. Frappe, A. Kaiser and J. Rabaey, "Capacitive body-coupled communication in the 400–500 MHz frequency band," in *EAI Conference on Body Area Networks*, October 2019.

[10] G. Tochou, et al., "A Fully-Digital 0.1-to-27 Mb/s ULV 450 MHz Transmitter with sub-100 µW Power Consumption for Body-Coupled Communication in 28 nm FD-SOI CMOS," in *RFIC*, June 2021.

[11] F. Solt, et al., "Energy Efficient Heartbeat-Based MAC Protocol for WBAN Employing Body Coupled Communication," *IEEE Access,* pp. 182966 - 182983, October 2020.

[12] A. Moin, A. Thielens, A. Araujo, A. Sangiovanni-Vincentelli and J. M. Rabaey, "Adaptive Body Area Networks Using Kinematics and Biosignals," *IEEE Journal of Biomedical and Health Informatics,* vol. 25, no. 3, pp. 623-633, 2020.

[13] H. De Man, "Ambient Intelligence: Giga-scale Dreams and Nano-scale Realities,," in *Keynote Presentation, ISSCC*, San Francisco, 2005.

[14] G. Templeton, "Smart Dust - A computer that is smaller than a grain of sand," May 2013. [Online]. Available: https://www.extremetech.com/extreme/155771-smart-dust-a-complete-computer-thats-smaller-than-a-grain-of-sand.

[15] D. Seo, et al., "Wireless Recording in the Peripheral Nervous System with Ultrasonic Neural Dust," *Neuron,* vol. 91, no. 3, pp. 529-539, 2016.

[16] P. Kanerva, "Hyperdimensional Computing: An Introduction to Computing in Ditributed Representation with High-Dimensional Random Vectors," *Cognitive Computing,* vol. 1, pp. 139-159, 2009.

[17] R. Kaveh and et al., "Wireless User-Generic Ear EEG," *IEEE Transactions on Biomedical Circuits and Systems,* vol. 14, no. 4, pp. 727-737, 2020.

[18] E. D'Anna, "How Kernel Plans to Put a Brain-Computer Interface in Every US Home by 2033," Neurotech Bites, November 2020. [Online]. https://edoardodanna.ch/article/kernel_flow_in_every_home_by_2033

[19] E. Waltz, "Elon Musk Announces Neuralink Advance Toward Syncing our Brains with AI," *IEEE Spectrum*, August 2020