Energy efficient heartbeat-based MAC protocol for WBAN employing body coupled communication

Flavien Solt, Robin Benarrouch, Guillaume Tochou, Oliver Facklam, Antoine Frappé, Andreia Cathelin, Andreas Kaiser, Jan Rabaey

Energy Efficient Heartbeat-Based MAC Protocol for WBAN Employing Body Coupled Communication

F. SOLT[1,2,3], R. BENARROUCH[2,3,4], (Student Member, IEEE), G. TOCHOU[2,3,4], (Student Member, IEEE), O. FACKLAM[1,2,3], A. FRAPPÉ[4], (Senior Member, IEEE), A. CATHELIN[2], (Senior Member, IEEE), A. KAISER[4], (Senior Member, IEEE), and J. RABAEY[3], (Fellow, IEEE)

[1]École Polytechnique, IP Paris, 91128 Palaiseau, France
[2]STMicroelectronics, 38920 Crolles, France
[3]BWRC, University of California Berkeley, Berkeley, CA 94704, USA
[4]Univ. Lille, CNRS, Centrale Lille, Univ. Polytechnique Hauts-de-France, Yncrea Hauts-de-France, UMR 8520 - IEMN, F-59000 Lille, France

Corresponding author: Flavien Solt (e-mail: flavien.solt@polytechnique.org).

ABSTRACT Wireless Body Area Networks (WBANs) are a fast-growing field fueled by the number of wearable devices developed for countless applications appearing on the market. To enable communication between a variety of those devices, the IEEE 802.15.6 standard was established. However, this standard has some intrinsic limitations in addressing the heterogeneity of the network nodes in terms of activity, data rates (from less than $\mathrm{bit/s}$ to multiple $\mathrm{Mbit/s}$), energy availability, form factor, and location on, around or inside the body. To address these concerns, an alternative model is proposed that could serve as an extension of the IEEE 802.15.6 Standard. At its core is an adaptive and low-overhead synchronization scheme based on heartbeat sensing. This forms the base for a TDMA-based (Time Division Multiple Access) Media Access Control (MAC) protocol dedicated to multi-tier networks. While this effort focuses specifically on Capacitive Body-Coupled Communication (C-BCC), other physical layers can be easily incorporated as well. Based on these premises, this paper compares various random-access slot allocation approaches to accommodate the multiple data rates matching the system requirements, while incorporating a duty-cycling strategy anchored by heartbeat detection. This work proposes a novel, flexible, and robust solution, making use of heartbeat synchronization and addressing the corresponding challenges. It efficiently interconnects multiple device types over a wide range of data rates and targets a mesh of stars topology. At the cost of an increased communication latency, the proposed protocol outperforms the IEEE 802.15.4 MAC standard in terms of energy efficiency by a factor of at least 12x in a realistic scenario.

INDEX TERMS Bio-signal, Body Area Network (BAN), Communication Protocol, Energy Efficiency, Medium Access Control (MAC), Synchronization, Wearable, Wireless Sensor Networks (WSN).

I. INTRODUCTION

Demand for medical and wellness applications grew extensively within the last decade, increasing the interest in Wireless Body Area Networks (WBANs). The most generic yet complex human-oriented concept, the Human Intranet [1], aims at interconnecting a wide variety of sensors and actuators. Such versatility relies on an extended ability to integrate various devices in terms of data rate, available power and communication latency among others.

WBANs have been extensively studied at the Media Access Control (MAC) level [2], [3]. Different protocols address trade-offs in terms of energy efficiency, latency, channel availability, and flexibility. Communication standards such as IEEE 802.15.6 [4], and, to some extent, IEEE 802.15.4 [5], have seen some penetration in WBAN deployment.

A number of physical layers over different frequency bands have proven to be effective (NB and UWB RF, inductive, capacitive, etc). In this paper, we pay special attention to sub-GHz Capacitive Body-Coupled Communication (C-BCC), a prime candidate for the physical layer of human-centered body communication, due to its efficiency and robustness. However, to reach high performance at a low energy cost, co-optimization between PHY and MAC layers is mandatory. A crucial challenge in reaching the best energy efficiency in managing communications in hybrid and versatile BANs supporting widely varying workloads such as the Human Intranet [1], is the establishment of a reliable synchronization scheme. While synchronization could be accomplished through the broadcast of beaconing signals, or rendezvous and wake-up schemes, the overhead of these

IEEE Access (ISSN : 2169-3536) 05-October-2020. DOI: 10.1109/ACCESS.2020.3028800

is substantial. An alternative approach could be through the detection of global body-wide bio-signals, such as the heartbeat, an approach that is still in its infancy but holds great promises. Using the heartbeat as a synchronization signal offers obvious advantages: it frees the system from generating a synchronization signal and each node from embedding a high accuracy timing block, usually bulky and power-consuming. Since the heart is the unique beacon generator, there are no synchronicity concerns other than its propagation delay. Moreover, state-of-the-art heartbeat detectors reach energy efficiencies that make this approach appealing in an ultra-low-power setting.

Heartbeat synchronization is at the core of the H-MAC protocol, introduced in [6], but this protocol limits its topology to a star network and trades off robustness and flexibility for device battery life extension. To the authors' knowledge, [6] and [7] were the only two attempts to use heartbeats as a synchronization signal until today. Reference [7] takes advantage of the heartbeat to synchronize and enable data exchange between different types of nodes. It uses the heartbeat as a central clock for all the nodes of the network. This method, if exploited properly, provides a reliable always present zero-energy (for the network) synchronization signal. If a node cannot efficiently sense the heartbeat, a local or WBAN-global lighthouse approach can be imagined with a node broadcasting the heartbeat signal. This topic is not covered in this paper.

In a related effort, [8] proposes a beacon-based synchronization MAC approach for the WBAN, which shares similarities with the present heartbeat synchronization capable context. It proposes a scheduled Time Division Multiple Access (TDMA) scheme, splitting frames into sensor-to-hub and hub-to-hub communication. However, the scheme targets two-tier star topologies, and the communication schedule is fixed. It assumes similar communication latency and volume for all sensor end-nodes around a given master node, which does not provide efficient support for various data rates and event-driven traffic.

The LDC-MAC protocol, introduced in [9], proposes a new beacon-based protocol that regroups essential features for a real WBAN: relatively low power consumption and effective cohabitation between low and high data rate traffics. To reduce energy consumption, LDC-MAC tends to use long superframes, in which regularly spaced slots are inserted to reduce the communication latency. However, a concurrent attempt to reduce power consumption in low-power communicating devices may lead to substantial clock drift [7], limiting practical superframe lengths. Heartbeat synchronization outperforms this trade-off.

Both IEEE 802.15.4 and 802.15.6 standards are not suitable for the targeted applications. Their non-beacon operation modes are mostly suited for low-traffic environments, not appropriate in our context [9]. On the other hand, their beacon operation mode suffers from the overhead addressed by LDC-MAC in [9]. Furthermore, both standards specify frequencies of operation and bandwidths which do not reach the Human Body Communication characteristics highlighted in [10].

Heartbeat synchronization presents some unique advantages as the synchronization beacon is globally available and not system generated. By nature, it addresses issues present in beacon-enabled protocols [11], such as the need for a central time coordinator or expensive distributed synchronization mechanisms. Indeed, heartbeat synchronization frees the network from a complex organization and selection process to define a node in charge of broadcasting the synchronization signal. This increases the system's robustness by not relying on a unique node in case of a failure. It also enables a completely distributed mesh operation, without the need for additional beacon schedule management. Furthermore, heartbeat-based synchronization does not suffer from any range limitation. Ultimately, no bandwidth allocation is necessary for broadcasting the synchronization signal, a desired feature while dealing with single-channel PHY such as C-BCC.

However, heartbeat-based synchronization comes with its specific challenges. The first is the heartbeat period's local and global variability. Local variability corresponds to the difference between two consecutive heartbeat interval durations. It is modeled in the present work as a normal distribution with a standard deviation of σ= 30 ms [12]. Global variability refers to the heart rate (HR) upper and lower bounds. It is assumed to fall in the range from HR_{min}=36 beats per minute (bpm) to HR_{max}=210 bpm [13] [14]. The second challenge is the non-negligible heartbeat signal propagation time. Its propagation speed is at least v_{HB}>250 m/s [15]. It requires device calibration in addition to substantial guard intervals. Ultimately no MAC protocol is today suitable to address the requirements of modern and future WBANs while benefiting from a heartbeat-based synchronization.

To accommodate the cohabitation between distinct classes of data rates, a TDMA protocol supporting a puncturing feature is adopted in [7]. The main high data rate traffic (i.e. between primary nodes) is periodically interrupted allowing secondary nodes to upload their data. The principle is illustrated with a simple 2-node example in Fig. 1, and forms the base for the TDMA scheme adopted in this paper.

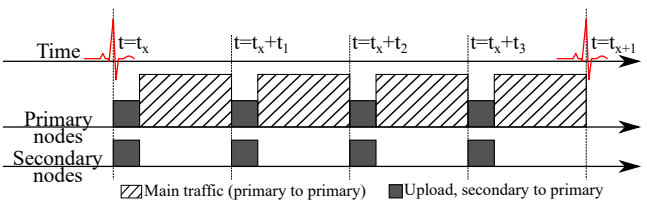

FIGURE 1. Main traffic periodically punctured.

The proposed work presents HB-MAC, a flexible and realistic MAC protocol for the Human Intranet. It effectively supports a wide span of high and low data rates as well as event-driven sensor traffic and emergency signaling. HB-MAC has been designed to address the intrinsic challenges of

heartbeat synchronization, taking advantage of this synchronization to lower the global energy consumption.

This paper is organized as follows: Section II presents the Human Intranet and physical layer requirements. Section III presents the proposed MAC protocol. Section IV presents metrics and simulation results. Section V discusses these results and compares them to prior art. Finally, Section VI concludes this work.

II. HUMAN INTRANET NETWORK REQUIREMENTS

A. THE APPLICATION REQUIREMENTS

The Human Intranet is meant to interface a wide variety of wearables. From [16], [17] and [18], the amount of data generated for different kinds of sensors (e.g. ECG, Temperature, motion...) is evaluated and synthesized in Table 1.

TABLE 1. 12-bit resolution sensor data generation rate

Signal	Sampling rate (Hz)	Data generation rate (bit/s)
ECG	120-250	1440-3000
Temperature	0.2-2	2.4-24
Oximetry	60	1440
Respiration rate	20	240
Heart rate	10	120
Biometric Z	10-20	120-240
Chemicals	10	120
Motion (/axis)	100-250	1200-3000
Neural recording (/channel)	10k - 30k	120k - 360k

From Table 1, the need for two types of nodes can be identified: those generating a limited amount of data, such as heart rate and temperature monitoring, and those generating a significant quantity such as motion sensing or neural recording which can stream around 25 $\mathrm{Mbit/s}$ for a 64-channel implant.

We define those two types of nodes as *hubs* and *leaves* [7]. The former offers greater computing capabilities and energy availability, whereas the latter is a lot more constrained and is energy-frugal.

In terms of behavior and role within the network, a leaf is connected to a unique hub, also called *parent hub*. A hub, however, can be connected to multiple leaves and hubs at the same time (see Fig. 2).

B. THE PHYSICAL LAYER

There are many options in developing a wireless network that spans the human body, each of which has its advantages and challenges. Often heterogeneous approaches combining multiple communication modalities are required. Considering the needs formulated in the paragraph above, one option that addresses both the upper and lower performance requirements is the use of Human-Body Coupled Communication in the 400 to 500 MHz band. While this choice is not essential, it helps in providing a unifying perspective.

In [10], it was established that a C-BCC solution operating in the 400-500 MHz band (with a center frequency of 450 MHz) offers some distinct advantages. At these

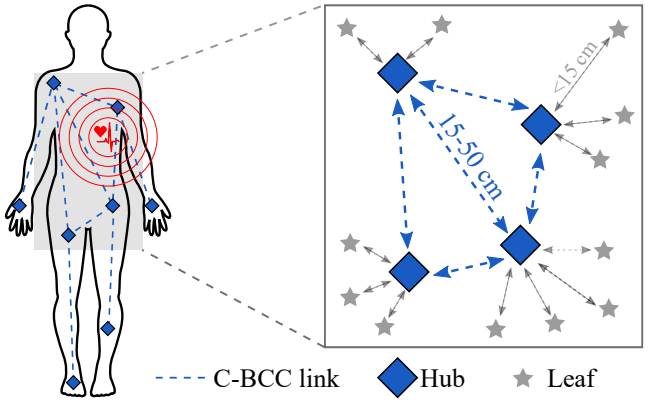

FIGURE 2. Human Intranet Network architecture.

frequencies, surface waves are the dominant propagation mechanism, making the channel less sensitive to changes in the environment and to coupling with external references. Its form factor is smaller and more convenient for the end-user than antenna-based RF technologies. Its attenuation per unit distance is also lower [19], at the cost of higher insertion loss. C-BCC matches the needs set by the applications and the developed scheme is tuned given C-BCC specifications. It is adaptable to other technologies, not covered in this paper.

From measurements [10], it was established that a bandwidth of at least 100 MHz is available with flat attenuation. The derived channel loss model is presented in (1), and its parameters as measured at 450 MHz are enumerated in Table 2.

$$L = \begin{cases} L_0 & d = d_0 \\ L_0 + \alpha_0(d - d_0) & d_0 < d < d_1 \\ L_1 + \alpha_1(d - d_1) & d_1 < d \end{cases} \quad (1)$$

TABLE 2. Channel loss model parameters

Frequency (MHz)	L_0 (dB)	α_0 (dB/cm)	L_1 (dB)	α_1 (dB/cm)	σ (dB)
450	-47	-0.95	-55	-0.37	6.5

σ is the standard deviation of the measurement compared to the theoretical fit, explained by the non-homogeneity of the medium (i.e. the human body). L_0 expresses the loss for the shortest distance recorded at $d_0 = 5\,\mathrm{cm}$. L_1 represents the loss at $d_1 = 15\,\mathrm{cm}$, transition distance between the quasi-static (QS) and surface wave regimes.

To take full advantage of such bandwidth, pulse-based communication is envisioned. It allows an aggressive transmitter duty-cycling leading to energy-efficient data transmission. Based on the analysis from [20], for a windowed-cosine pulse model, a 20 ns Gaussian pulse in the time domain corresponds to a 150 MHz-bandwidth signal at $-10\,\mathrm{dBc}$ in the frequency domain.

With a simple modulation such as OOK or BPSK transmitting 20 ns pulses, the maximum theoretically achievable data rate is 50 $\mathrm{Mbit/s}$.

Given a similar architecture as presented in Fig. 2, the type and amount of data exchanged from Table 1, the channel characterization results from [10], and the pulse-based communication centered on 450 MHz, the hub-to-hub communication data rate is set to 50 $\mathrm{Mbit/s}$ with BPSK modulation scheme. The leaf-to-hub link is however set to 100 $\mathrm{kbit/s}$ with OOK modulation scheme. It offers a good trade-off between the transmission time of a given message and the achievable transmitter energy efficiency, relaxing the constraints on the leaf timing.

The receiver sensitivity S is calculated in (2) as a function of the integrated thermal noise N_{th} over the bandwidth of consideration BW (see (4)) [21].

$$S = N_{th} + NF + SNR_{req} \tag{2}$$

where NF is the receiver noise figure and SNR_{req} is the required signal-to-noise ratio expressed in (3) as a function of the desired data rate DR.

$$SNR_{req} = \frac{E_b}{N_0} + 10 \cdot log_{10}\left(\frac{DR}{BW}\right) \tag{3}$$

$$N_{th} = 10 \cdot log_{10}(K \cdot T \cdot 10^3) + 10 \cdot log_{10}(BW) \tag{4}$$

For a usual noise figure of $NF = 6$ dB, a bandwidth of $BW = 150$ MHz, a ratio $E_b/N_0 = 14$ dB and a data rate of $DR = 100$ $\mathrm{kbit/s}$ and $DR = 50$ $\mathrm{Mbit/s}$, the sensitivity becomes respectively:

$$S_{100\mathrm{kbit/s}} = -104\mathrm{dBm} \;\&\; S_{50\mathrm{Mbit/s}} = -77\mathrm{dBm}$$

In terms of geometry, the leaf-hub distance is limited to 15 cm, which means from equation 1, that the signal attenuation is not greater than 55 dB. Given a receiver sensitivity of -104 dBm and a 3 dB margin, the minimal output power the leaf transmitter should radiate is about -46 dBm. The inter-hub range is set to 50 cm (about an adult arm length) allowing a large coverage of a human body. Applying the same calculation, the hub transmitter output power should be around -10 dBm.

The above-listed network details regarding the communication mechanism, distance covered, channel attenuation, signal modulation, and transmitted power set the constraints affecting the HB-MAC protocol.

III. COMMUNICATION & MAC LAYER PROPOSAL

A. MODEL & ARCHITECTURE

The proposed MAC scheme, HB-MAC, is a TDMA-based scheme adapted for two-stage networks. Each hub surrounded by its leaves forms a structure called a *cluster*, while the hubs are connected in a mesh. The considered network topology is thus, a mesh of stars. Given the substantial fading over distance that comes with Body Coupled Communication, it is assumed that intra-cluster communications (i.e., between a hub and a leaf) do not interfere with intra-cluster communications in other clusters. Thus, at any time, all leaves in the network have at most one *parent hub*.

While clusters are formed by positioning leaves around hubs, HB-MAC is designed to be adaptable. The link is dynamically adjusted to mitigate for body dynamics. This approach brings system flexibility and robustness against connection/disconnection. It also offers the ability to add, move or remove a leaf from the network during runtime without impacting the other devices.

Leaves can interact with their parent hub in two different ways:

- In ***detached mode***, a leaf communicates opportunistically upon the occurrence of an event. Detached communication can also be used to initiate an attachment procedure.
- In ***attached mode***, a leaf is connected to its hub for a longer term. It reduces the overhead by avoiding random-access procedures. Moreover, an attached leaf has a well-defined location in the network and can consequently be addressed more effectively by distant hubs.

It is further assumed that all devices embed heartbeat sensing capability and the relative detection time offset has been calibrated. This offset is caused by a non-negligible heartbeat signal propagation time [15].

HB-MAC uses *superframes* as its main structure. Superframes are defined as the time interval between two heartbeats. Since the signal skew is dependent on the device location on the body, guard intervals are introduced. The present protocol uses two different forms of superframes:

- The *regular superframe* is the most represented and is shown in Fig. 3. This superframe type only allows attached and inter-hub traffics, using a protocol based on existing standards for the latter.
- The *detached superframe* allocates time to both detached leaves and hub-to-hub communication as shown in Fig. 4. Detached superframes are less frequent than regular superframes, as their purpose is to serve rare or event-driven traffic, including emergency communication and management procedures.

A detached superframe may appear for two different reasons. It occurs at pre-defined rates corresponding to typical leaf latency requirements. The second possibility is in case of an emergency. When a leaf calls for an emergency (as described in Section III-B), a detached superframe is started both for the concerned hub and its direct neighbors. In this second case, the detached superframe is referred to as an *emergency superframe*.

Except in the emergency case (either triggered locally by a leaf using a dedicated slot or initiated by a hub), detached superframes occur simultaneously over the whole body. The detached superframes' phase and period of occurrence are agreed by the hubs. This could be accomplished by any distributed algorithm, and in the most straightforward case simply decided by a mother hub. This scheme allows for dynamic reallocation in case of a workload change.

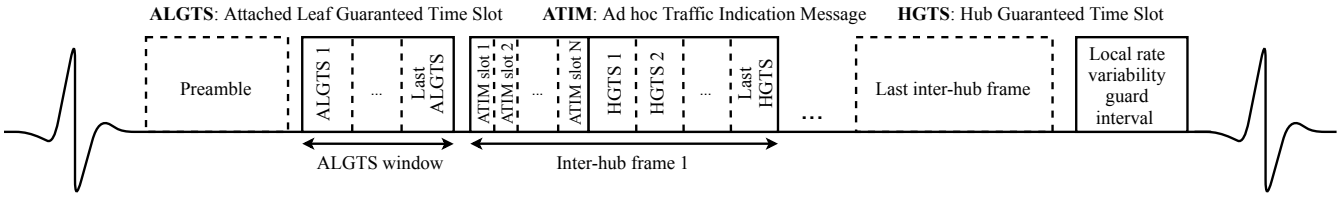

FIGURE 3. Regular superframe structure.

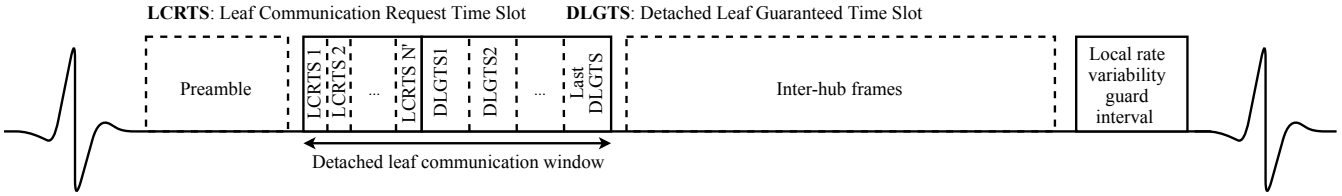

FIGURE 4. Detached superframe structure.

B. SUPERFRAME PREAMBLE

All superframes, regardless of their type, start with the same preamble scheme. The preamble's purpose is twofold. First, it provides an interface for low-latency emergency signaling. Second, it allows leaves to know when the next detached superframe will occur. Its structure is illustrated in Fig. 5.

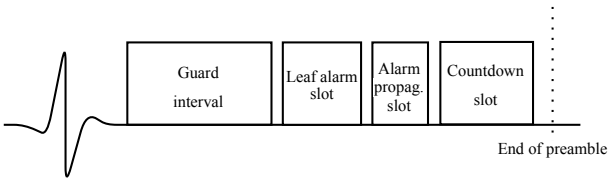

FIGURE 5. Superframe preamble structure.

1) Leaf Alarm Slot

It is crucial for a WBAN MAC protocol to support low-latency emergency signaling, allowing data upload with minimal delay. This must have minimal impact on the regular operation of the network.

It is implemented by reserving an emergency slot immediately after the post-heartbeat guard interval. This is an analog slot: if a hub senses energy, then it enters the emergency mode. Precisely, the latter hub propagates a specific signal, sets its countdown value (described below) to the emergency value, and enters an emergency superframe.

This emergency scheme is robust against cardiac arrest, as any hub is capable of sensing the heartbeat. If the elapsed time from the last detected heartbeat exceeds a predefined limit, it reacts accordingly.

2) Alarm Propagation Slot

An emergency state cannot be strictly local to a given cluster. First, an emergency event is likely to be a global event happening within the human body, hence is likely to be triggered in different clusters simultaneously. Second, all neighboring hubs must be aware of the emergency mode to avoid interfering with the unpredicted communication between the hub in alert and its leaves.

An alarm propagation slot is thus implemented after the alarm trigger slot, which is intended for hub-to-hub emergency noticing. A hub sensing a signal during the alarm propagation slot enters an emergency superframe as well.

3) Leaf Countdown Slot

Not requiring any attachment procedure is a beneficial feature when targeting ultra-low power consumption on the leaf side. This is enabled by the introduction of a countdown slot. It also gives leaves the ability to detect a parent hub change.

Periodically, hubs broadcast a countdown value to their leaves. The countdown message carries the hub address along with the number of heartbeats separating the current superframe from the next detached superframe. Transmitting the hub address ensures that a leaf communicates with the intended hub, especially in attached mode. The countdown value is typically encoded as an 8-bit unsigned integer, since the highest representable value corresponds to at least 73 seconds, for any heart rate in the considered range. According to the typical data generation rates given in Table 1, this upper-bound duration is sufficient. A specific countdown value indicates that the current superframe is in an emergency state. The countdown slot can be extended to contain additional information, such as the random-access strategy to adopt by detached leaves for instance.

The countdown system allows a leaf to hibernate for an indefinite amount of time, turning off communication and

heartbeat sensor, and to opportunistically enter the network without other overhead. This is of particular interest not only for long hibernations but also to minimize the overhead of leaf re-attachments due to changes of the network topology, typical for Human Body Communications.

Finally, the countdown slot is also an acknowledgment informing a potential alarm-triggering leaf whether the alarm signal has been received and accepted by the hub.

C. REGULAR SUPERFRAME

1) General Structure

The regular superframe is divided into two parts as depicted in Fig. 3. The first part is dedicated to communication with attached leaves during the Attached Leaf Guaranteed Time Slots (ALGTS), whose internal structure is application dependent.

The second part is dedicated to inter-hub communication. Its precise specification is not crucial to the present work. We propose to divide this part into multiple smaller frames similar to the 802.11 PSM protocol [22]. Each frame includes a reservation sub-window following a CSMA/CA scheme, materialized by the use of ATIM (Ad hoc Traffic Indication Messages) time slots, and then a HGTS (Hub Guaranteed Time Slots) sub-window. This structure has several advantages. First, it avoids the idle listening overhead from random-access schemes such as simple or slotted Aloha. Second, keeping the frames short improves both throughput and reliability. The smaller the frames, the faster the packets can flow across the entire network since a packet cannot perform more than one hop per frame. In addition, longer frames are naturally more prone to be affected by channel changes, common for body deployed networks [23]. Finally, as shown further below, collisions between concurrent communications are inevitable and may become frequent in some configurations. The choice of a division between reservation and guaranteed slots reduces the collision's energy cost [24].

2) Heartbeat Variability Management

One important parameter impacting the scheme efficiency is the length of the guard interval at the end of the superframe illustrated in Fig. 6. It is required to compensate for the heart rate local variability. The standard deviation between two successive heartbeat signals is 30 ms. The heartbeat occurrence can be predicted based on the previous superframe. Overestimating the available time can lead to packet failure and increased energy consumption.

3) Re-scheduling procedure

Leaves should be reschedulable. During a hub's ALGTS window, neighboring hubs are not allowed to communicate with other hubs. Thus, it is beneficial to locally re-schedule leaves to minimize the hub idle duration due to neighboring hubs still communicating in their ALGTS windows.

The rescheduling algorithm requires acknowledgments from the reallocated leaves. The time spent waiting for all the required acknowledgments through successful ALGTS

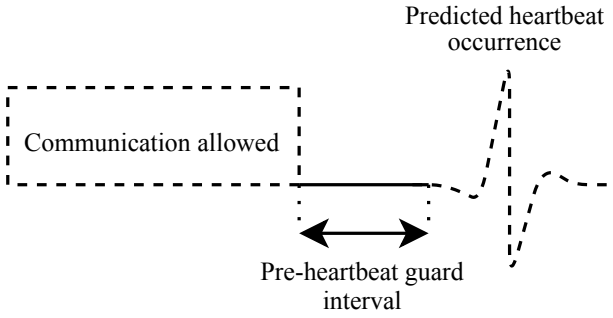

FIGURE 6. Pre-heartbeat guard interval.

is referred to as the *rescheduling latency* and can cause a global leaf rescheduling procedure to last long. In the worst case, hubs need to wait for all the leaves' attachments to time out. Moreover, rescheduling requires the use of additional temporary ALGTS slots.

D. DETACHED SUPERFRAME

1) General Structure

Only detached superframes allow detached leaves (i.e. leaves which haven't completed a still valid attachment process) to communicate. They allow leaves to communicate with the hub in range and then let the hubs communicate with each other for the remaining duration of the superframe. The uplink traffic is privileged. To receive downlink data, a leaf has to initiate the communication with the hub, which is a common decision among energy-efficient protocols, such as LoRaWAN Class A [25]. The detached superframe structure shares similarities with the regular superframe structure. It substitutes the ALGTS window with a detached leaf communication window, illustrated in Fig. 4.

2) Attachment procedure

Only an overview of the attachment mechanisms is provided. It underlines the challenges faced by such procedures. The attachment procedure is performed during a detached superframe. Being attached provides two advantages: (i) facilitated communication with the hub, and (ii) a fixed address in the network, allowing data to be routed to the leaf through the WBAN.

To get attached, a leaf sends to the hub the preferred period of communication measured in superframes, its communication class which defines the ALGTS internal structure (e.g., actuator-type communication or sensor-type communication), data related to the communication class (e.g. its data generation rate for a sensor-type communication) and finally a list of identifiers dealing with the types of data to be addressed to the leaf by other devices in the WBAN. The hub immediately responds with information regarding the scheduled phase and period, as well as the dedicated ALGTS.

3) Detached Leaf Communication Window

A detached leaf communication window is illustrated in Fig. 4 as part of the detached superframe. It is structured with an initial CSMA/CA sub-window, followed by a sub-window of guaranteed time slots. The first sub-window contains Leaf Communication Request Time Slots (LCRTS), each followed by an acknowledgment slot for the hub. These Leaf Communication Requests (LCR) sole purpose is to reserve the following Detached Leaf Guaranteed Time Slots (DLGTS), which are followed by an acknowledgment slot.

The size of both LCRTS and DLGTS sub-windows is later addressed in this paper. It depends on the number of leaves in a cluster and on the target leaf energy consumption. The latter trades off with the time remaining for inter-hub communication. Typically, in an emergency superframe, the dimensions of the detached leaf communication window are larger than during regular operation to maximize the probability of communication success, as the leaf uplink traffic may brutally increase.

Locating leaf-to-hub communication at the beginning of the detached superframe offers several advantages. First, it increases the likelihood that the leaves can still communicate with the hub, once they have read the zero-countdown value (described in Section III-B) since WBAN does not guarantee a constant channel quality between two fixed devices on the body [23], [26]. It also serves as a detached superframe occurrence confirmation and Quality of Service (QoS) information. Second, leaves may be equipped with inaccurate clocks [7]. Hence, communicating early in the detached superframe avoids excessive clock drift accumulation since the clock has been reset by the heartbeat detection. Finally, this detached superframe arrangement allows information to be spread as deep as possible in the network, and thus offers relatively low end-to-end latency.

4) Post Detached Leaf Communication Window

Once the leaf communication window is complete, the scheme proceeds to inter-hub communication. Typically, no ALGTS is inserted behind the detached leaf communication window for two reasons. First, the length of the detached leaf communication window is not fixed. It is allowed to temporally vary to accommodate the workload, as long as it has the same global value spatially across the network. Second, scheduling a consecutive attached communication window without a substantial risk of interruption by the next heartbeat would be too constraining.

It is not always satisfying to forbid collisions between leaves granted ALGTS (in terms of phase and period of ALGTS occurrence) and detached superframes. The detached window period may vary, as long as it remains global to the WBAN. Leaves could be rescheduled accordingly, but rescheduling latency could substantially delay the modifications of the detached communication period. Ultimately, enforcing that no ALGTS ever coincides with a detached superframe, is equivalent to ensuring (5):

$$gcd(T_A, T_D) \nmid \phi_A \tag{5}$$

where T_A represents the scheduled ALGTS period, ϕ_A the scheduled phase with reference to any occurrence of a detached superframe, and T_D the detached superframe period measured in heartbeat intervals.

According to the value of T_D, this condition may become restrictive, as it forbids, for instance, T_D to be prime with T_A. This may lead to choosing detached superframe periods accordingly.

If a collision occurs, either because the condition given in (5) has knowingly been infringed or because of an emergency event, leaves can either wait for their next ALGTS, or communicate with the hub through the detached communication window, at the cost of a random-access. Such collisions between ALGTS-containing superframes and detached superframes are dynamically detected by attached leaves, as they read the countdown value right before proceeding to attached guaranteed communication.

IV. METRICS AND SIMULATION RESULTS

Within this section, the HB-MAC protocol performance is evaluated. We focus on the detached mode as this is where the innovation is: the attached leaf operation and the inter-hub communication rely on existing protocols. First, the dimensions of the leaf reservation window will be analyzed and optimized in terms of energy cost and communication failures. The average detached leaves' ON-time is derived, and the impact on the hubs. Finally, a heartbeat detector duty cycling strategy, proper to this MAC protocol, is proposed to further improve the leaf energy efficiency.

The protocol parameters used in the analysis are given in Table 3 and Table 4. The overhead of the physical header consists of a preamble, a frame delimiter, and a CRC of sizes 7, 1, and 4 bytes, respectively. Leaf and hub addresses have sizes of 32 and 16 bits respectively. Guard intervals of 1ms are used for communications with leaves, sufficient to mitigate their clock inaccuracy over one heartbeat interval [7]. The remaining values are calculated based on the data rates and packet structures.

TABLE 3. Elementary field sizes

	Overhead on PHY	Leaf address	Hub address	Countdown incl. addr.
Field size (bit)	96	32	16	24

TABLE 4. Slot size and duration, incl. guard intervals

	Countdown	Leaf com. REQ+ACK	Leaf GTS+ACK
Size (bit)	120	144	248 + data
Duration (ms)	2.20	4.54	4.6 + 0.0103·data

A. DETACHED COMMUNICATION STRUCTURE

1) Challenges and definitions

The detached leaf communication window dimensions result from a trade-off between the leaf communication success probability, hub energy consumption and inter-hub communication time. The number of DLGTS per detached leaf communication window should match the predicted number of leaves whose LCR succeeds. The corresponding negotiation window size calculation, on the other hand, requires attention. The LCR communication success rate, from a detached leaf to its parent hub, increases with the number of LCRTS available per superframe. However, an oversized LCR sub-window substitutes hub sleep or inter-hub communication time.

The number of leaves competing for simultaneous access to a specific hub, as well as the LCR random allocation strategy, significantly impact the negotiation success rate and energy efficiency. Given the energy scarcity, a leaf-dedicated power management strategy is required.

We define the *failure rate* (and its *success rate* counterpart) as the probability of unsuccessful leaf-to-hub communication, within a given LCR sub-window. Besides the physical channel variations or interference with the environment, the major source of transmission failures is the successive collisions during the LCR slots. It is the only element considered in this study. As there are no more than 80 bits sent per leaf per LCR, excluding the physical-layer preamble, and assuming a bit error rate of $10^{-6}/\mathrm{bit}$ [27], the failure probability of a LCR message due to physical deterioration is lower than $8 \cdot 10^{-5}$.

In the following, several probabilistic LCRTS allocation and back-off strategies are studied in terms of success rate and energy consumption, from a detached leaf standpoint. The energy consumption dedicated to random-access increases proportionally with the number of LCR messages sent by the leaf in a detached superframe.

This study focuses on *symmetric* strategies: all detached leaves trying to concurrently access the same hub follow the same rules. Since the set of symmetric LCRTS allocation and back-off strategies is large and difficult to parametrize, we only consider a set of common slot allocation families.

1) ***Uniform without back-off (UBS):*** Each leaf chooses uniformly and independently the LCRTS to send a communication request, within the entire LCR sub-window. If it fails, no back-off is performed. This strategy does not take any parameter.
2) ***Uniform with back-off (UB):*** Similar to UBS, except in case of a communication failure. In this case, the leaf individually re-applies the same strategy on the remaining LCRTS in the current LCR sub-window. This strategy does not take any parameter.
3) ***Fixed contention window size (FCS):*** the LCR sub-window is divided into contention windows of equal size. In each contention window, the leaves that have not successfully sent an LCR message yet, select uniformly one LCR slot within the contention window. This family of strategies takes one parameter: the contention window size.
4) ***Binary exponential back-off (BEB):*** Similar to FCS, except regarding the contention window size. It doubles at each back-off and is bounded by a given value. This family of strategies takes two parameters: the initial and the maximal contention window size.

Both families, FCS and BEB, can be implemented following two behaviors, illustrated in Fig. 7:

1) *Exceeding bound (EB):* If the last contention window exceeds the LCR sub-window, then the remaining leaves still willing to communicate choose their slot within the entire contention window (instead of restricting to the remaining LCR slots). If a leaf chooses a slot beyond the last LCRTS, then it cancels its communication attempt and typically waits for the next detached superframe.
2) *Containing bound (CB):* If the last contention window goes beyond the LCR sub-window, then the remaining leaves attempting to communicate choose their slot only in the remaining LCRTS.

UBS is a particular FCS parametrization, which is itself a particular BEB parametrization. In the following they are considered distinct to underline their specificities and simplify the strategy when possible.

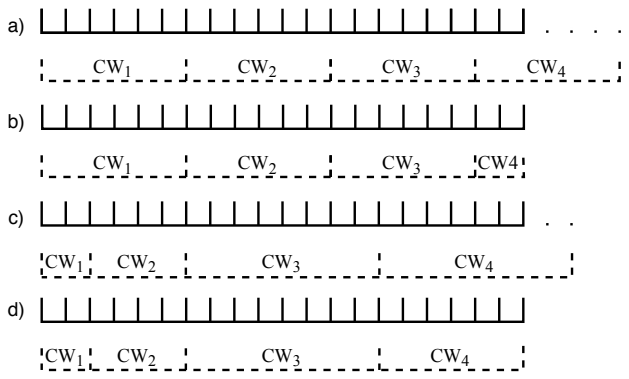

FIGURE 7. Successive back-off windows for a) FCS-EB with contention window size 6, b) FCS-CB with contention window size 6, c) BEB-EB with contention window bounds (2, 8), d) BEB-CB with contention window bounds (2, 8).

2) Performance evaluation

The various options are mathematically evaluated here. In the following formulae, the integer F represents the number of concurrent detached leaves trying to access the considered hub. The integer S represents the number of slots in the LCR sub-window, and M is the expected number of LCR messages sent by an involved detached leaf during the LCR sub-window. Since all considered strategies are symmetric among leaves, we can consider a fixed leaf F1 without loss of generality.

UBS: The leaf F1 succeeds its LCR communication if and only if all other concurrent leaves attempt communication

on a different LCR slot than the one chosen by F1. Equations (6a) and (6b) express the probability of success $P_{s,UBS}$ and message number M_{UBS} under this allocation scheme.

$$P_{s,UBS}(F,S) = \left(\frac{S-1}{S}\right)^{F-1} \tag{6a}$$

$$M_{UBS}(F,S) = 1 \tag{6b}$$

UB: The success probability is recursively calculated considering a partition of five cases happening on the first LCR slot, as indicated in Fig. 8. The evaluated number of messages sent uses the same event set partition and the observation that F1 sends a new message in the cases (i) and (ii) and not in the other cases. Using these partitions, we obtain the success probability $P_{s,UB}$ and expected number of messages M_{UB} sent by a candidate leaf during a LCR window given by (7a) and (7b) respectively.

$$P_{s,UB}(F,S) = \begin{cases} 1_{\{S>0\}} & \text{if } F=1 \\ 0 & \text{else if } S=0 \\ \frac{1}{S}\left(\frac{S-1}{S}\right)^{F-1} & \\ +\frac{1}{S}\left(1-\left(\frac{S-1}{S}\right)^{F-1}\right) P_{s,UB}(F,S-1) & \\ +\left(\frac{S-1}{S}\right)^{F} P_{s,UB}(F,S-1) & \\ +\left(\frac{S-1}{S}-\left(\frac{F+S-2}{S}\right)\left(\frac{S-1}{S}\right)^{F-1}\right) P_{s,UB}(F,S-1) & \\ +\left(\frac{S-1}{S}\right)^{F-1}\frac{F-1}{S} P_{s,UB}(F-1,S-1) & \text{else.} \end{cases} \tag{7a}$$

$$M_{UB}(F,S) = \begin{cases} 1_{\{S>0\}} & \text{if } F=1 \\ 0 & \text{else if } S=0 \\ \frac{1}{S}\left(\frac{S-1}{S}\right)^{F-1} & \\ +\frac{1}{S}\left(1-\left(\frac{S-1}{S}\right)^{F-1}\right)(1+M_{UB}(F,S-1)) & \\ +\left(\frac{S-1}{S}\right)^{F} M_{UB}(F,S-1) & \\ +\left(\frac{S-1}{S}-\left(\frac{F+S-2}{S}\right)\left(\frac{S-1}{S}\right)^{F-1}\right) M_{UB}(F,S-1) & \\ +\left(\frac{S-1}{S}\right)^{F-1}\frac{F-1}{S} M_{UB}(F-1,S-1) & \text{else.} \end{cases} \tag{7b}$$

BEB: The first step in computing the success probability and expected number of messages per concurrent leaf per LCR frame is to determine the number of applications $\phi : [1..F] \to [1..S]$ which verify the property $H(\phi)$: *The preimage of each element of the codomain of ϕ is either the empty set or a (finite) set of cardinal at least 2*. The number $Conf$ of such applications is calculated in (8). This serves as a tool to evaluate the number of LCR slots with unsuccessful LCR communication.

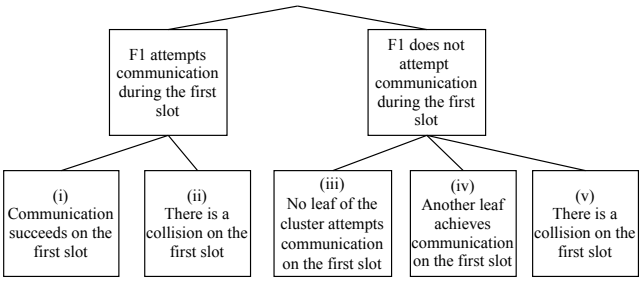

FIGURE 8. Case-based reasoning used in the recursive calculation in the UB strategy.

$$Conf(F,S) = \begin{cases} 1_{\{F\leq 0\}} & \text{if } S \leq 0 \\ 1 & \text{else if } F \leq 0 \\ Conf(F,S-1) + \sum_{n=2}^{F}\binom{F}{n} Conf(F-n,S-1) & \text{else.} \end{cases} \tag{8}$$

The second step counts the number of configurations $Conf_k$ (i.e., allocations ϕ of LCR slot indices to leaves) such that exactly k leaves achieve successful LCR communication with the hub. The number of possible configurations is calculated in (9).

$$Conf_k(F,S) = \binom{F}{k}\binom{S}{k} k! Conf(F-k,S-k) \tag{9}$$

From there comes the probability $P_{k,f}$ expressed in (10) that k concurrent leaves achieve successful LCR communication with the hub during the given LCR window and that F1's LCR communication fails.

$$P_{k,f}(F,S) = \frac{S\sum_{n=1}^{F-1}\binom{F-1}{n} Conf_k(F-n-1,S-1)}{S^F} \tag{10}$$

At this point, a partition is conditioned on the result of the first contention window. It allows to recursively compute the success probabilities $P_{s,BEB}$ in (11) and expected message numbers M_{BEB} in (12) for the four remaining schemes. The considered partition is the event set: {F1 succeeds during its first contention window} $\bigsqcup_{k=0}^{F-1}$ {F1 fails during its first contention window and k concurrent leaves succeed during this contention window}.

$$P_{s,BEB}(F,S,cw,cw_{max}) = \begin{cases} 0 & \text{if } S=0 \\ \frac{S}{cw}\left(\frac{cw-1}{cw}\right)^{F-1} & \text{else if EB and } S \leq cw \\ \left(\frac{S-1}{S}\right)^{F-1} & \text{else if CB and } S \leq cw \\ \left(\frac{cw-1}{cw}\right)^{F-1} + \sum_{n=0}^{F-1} P_{n,f}(F,cw) & \\ \cdot P_{s,BEB}(F-n,S-cw,\eta,cw_{max}) & \text{else.} \end{cases} \tag{11}$$

where:

$\eta := min(2cw, cw_{max})$

$$M_{BEB}(F,S,cw,cw_{max}) = \begin{cases} 0 & \text{if } S=0 \\ \frac{S}{cw} & \text{else if EB and } S \le cw \\ 1 & \text{else if CB and } S \le cw \\ 1+\sum_{n=0}^{F-1} P_{n,f}(F,cw) & \\ \cdot M_{BEB}(F-n,S-cw,\eta,cw_{max}) & \text{else.} \end{cases} \quad (12)$$

3) Results

HB-MAC aims at efficiently supporting detached communication. Focus is put on transmission reliability and energy cost. To pursue this objective, optimal strategies have been derived. Additionally, robustness also needs to be ensured, as the number of detached leaves attempting to access a given hub is difficult to anticipate.

For readability purposes, we systematically consider the optimal parameters for each family of strategies in this section. The decreasing priority has been set as UBS > UB > FCS/CB > BEB/CB > FCS/EB > BEB/EB. If several strategies offer identical performance, only the highest-priority strategy is given.

Derived from the previous analyses, the optimal LCRTS allocation and back-off strategies are illustrated in Fig. 9 for a minimal failure rate. This optimization provides a lower bound for the achievable failure rate due to collisions, along with the minimal expected number of messages per leaf to achieve it, as illustrated in Fig. 10. In particular, we observed that for nine leaves trying to concurrently access the same hub, the expected number of LCR messages sent in the LCR sub-window should never exceed 3.3 per concurrent leaf.

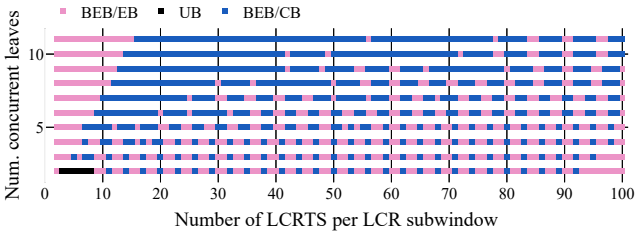

FIGURE 9. Optimal strategies to minimize the failure rate.

While the previous consideration focused on the failure rate as the optimization objective, another approach consists of optimizing the number of messages sent by a detached leaf under failure probability constraint. This approach is more realistic, since the HB-MAC aims at ensuring low energy consumption, whereas minimizing collision probability significantly below the physical transmission failure rate is practically pointless. Fig. 11 illustrates optimization results minimizing the expected message number under a failure rate constraint of 1%. Backed with more detail by the quantitative analysis, it shows the robustness of the optimization: optimal families of strategies remain stable for small variations in the number of LCRTS or concurrent detached leaves. The analysis proves that their parameters vary smoothly as well.

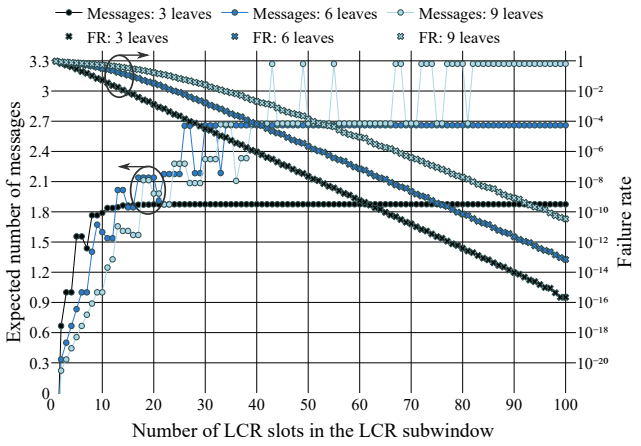

FIGURE 10. Minimal failure rate, and the corresponding expected number of messages sent, among the considered strategies.

Fig. 12 shows the actual minimized number of messages under the same failure rate constraint. Optimal spots appear, located at the slopes' inflection points. As their location is related to the number of concurrent leaves, they should be considered with the proper margin to propose efficient LCR window sizings.

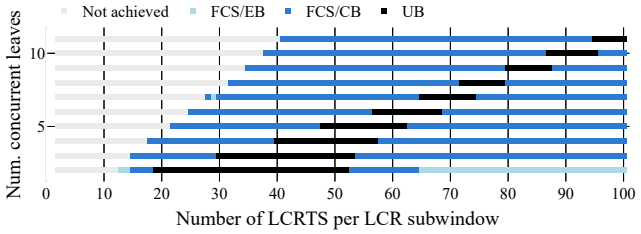

FIGURE 11. Optimal strategies to minimize the number of messages, under the constraint of failure rate lower than 1%.

As a numerical example, if a Human Intranet application around a given hub requires a failure rate below 1% and if 3 concurrent leaves try to access the hub during the same detached superframe, Fig. 12 shows that a LCR subwindow size of 30 LCRTS, reaches such an optimal point. Fig. 11 shows that the corresponding strategy is a uniform allocation with back-off.

The detached leaves' energy consumption is approximately proportional to the number of messages sent per LCR sub-window. In many cases, the choice of a sufficient number of LCR slots brings the message number below 1.2 per concurrent leaf even for dense clusters containing as many as nine detached leaves.

Those results show not only efficiency in terms of failure rate and energy consumption, but also the robustness of the detached operation. First, we demonstrated the existence of optimal slot allocation and back-off strategies that enable efficient detached operation for various numbers of LCR slots and for large numbers of concurrent leaves. Second,

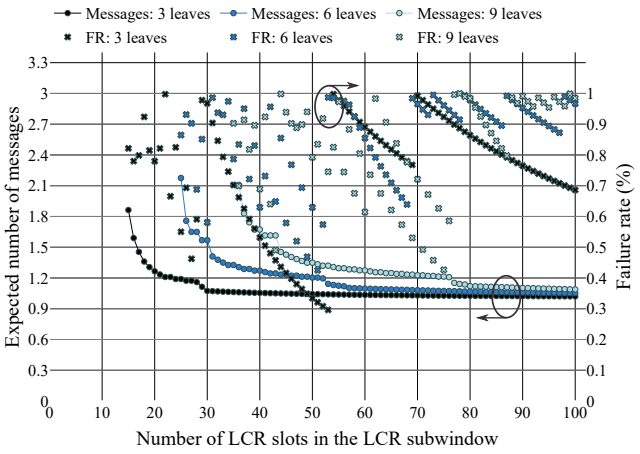

FIGURE 12. Minimal expected number of messages sent to achieve 1% failure rate, among the considered strategies.

we showed that the optimal strategies and their parameters vary smoothly. This smoothness is crucial to ensure that optimal strategies for a given predicted number of concurrent detached leaves remain efficient as the actual number varies.

Ultimately, one major objective of leaf attachment is to reduce the load on detached leaf communication windows. Only leaves with rare or irregular traffic, as well as leaves initiating attachment procedures and attached leaves requiring extra communication, should communicate in detached mode. This brings further down the actual expected number of LCR messages to send and the failure probability.

B. TRANSCEIVER POWER CONSUMPTION

1) Detached Leaf Management Overhead

Detached leaf integration costs time and energy. It impairs inter-hub communication by requiring dedicated time slots: superframe preambles and detached leaf communication windows. Extra hub energy is spent among four actions: (i) listening during the alarm propagation slot; (ii) sending periodic countdown information; (iii) listening during the entire LCR subwindow and, (iv) communicating during the DLGTS windows.

The average hub communication time overhead due to detached operation is highly dependent on the detached superframe period and is represented in Fig. 13, considering the numerical values listed in Table 4 and Table 5. The communication time overhead due to distributed operations (such as modifying the detached superframe period or advertising LCR subwindow sizes for modifications) is not covered in the study. It is worth noticing the large contribution of the LCR windows, estimated in the worst-case scenario (where the hub sends an acknowledgment at each LCR slot). The important contribution of LCR slots makes the hub time overhead sensitive to variations in size of the LCR window. Adapting the LCR window length to the expected detached traffic will have a substantial effect by allowing an efficient trade-off with hub communication time overhead dedicated to detached leaf management.

TABLE 5. Chosen figures for numerical considerations

Leaf data generation rate	200 bit/s
Countdown reads per detached superframe occurrence	2
Number of LCR slots	30
Device address bits	32 bits
Total physical overhead bits (preamble, SFD, and CRC)	96 bits
Transceiver power consumption at 100 kbit/s	100 µW [28]
LGTS payload	4096 bits

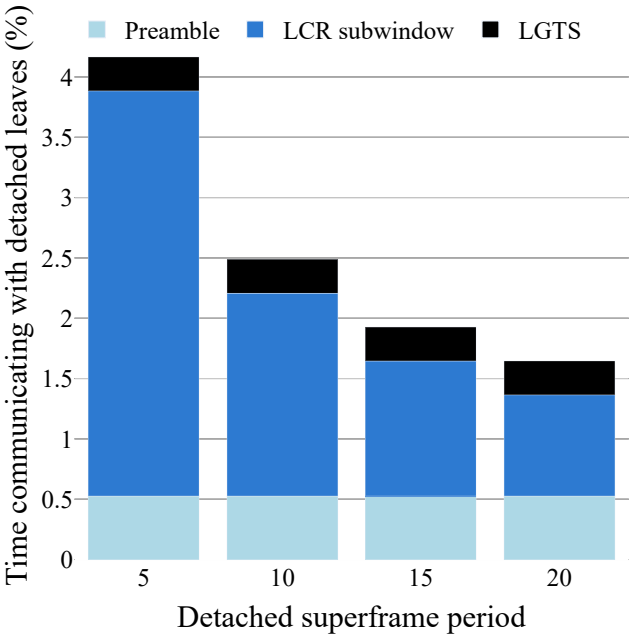

FIGURE 13. Percentage of the time dedicated by hubs to communication with detached leaves.

2) Detached Leaf Internal Overhead

A major strength of HB-MAC is the detached leaves' low energy consumption. It is brought to a very low value by minimizing the communication time. Detached leaves only communicate to (i) read countdown information; (ii) attempt LCR requests and listen to the corresponding acknowledgment, and (iii) during the corresponding DLGTS.

Even if the leaves can read the countdown slot only once per detached superframe while maintaining a functional structure, a second countdown read operation is considered. The second countdown read offers strong benefits such as channel quality assessment, ascertaining that the hub is currently reachable. The optimizations performed in Section IV-A resulted in a small expected number of LCR messages. Calculations were made for 1.1 LCR message sent, according to Fig. 12 for three concurrent detached leaves. The resulting average leaf ON-time is represented in Fig. 14. As most of the LCR communication overhead has been

moved to the hub, it takes a minor part of the total detached communication time. Moreover, the countdown and LCR overhead due to countdown reading and LCR communication can be brought to arbitrarily low values by tolerating higher communication latencies for a given detached leaf, granting very low energy consumption for sparsely communicating leaves.

Additionally, Fig. 14 shows the overhead of a detached operation compared to an ideal attached operation for different detached traffic sparsities. For instance, a leaf communicating every 20 superframes and using guaranteed time slot payloads of $4\,\mathrm{kbit}$, in concordance with realistic data rates summarized in Table 1, would reduce its communication time by 9.10% by running an attached mode. Attachment would additionally reduce the load on reservation slots.

Different leaf communication requirements are addressed by attached and detached modes. While leaving the precise attached operation aside in this study, HB-MAC introduces an efficient detached operation dedicated to event-driven and rare traffic. The protocol aims at keeping the detached overhead low, focusing on leaves' energy consumption, and we showed that it can efficiently be adapted to accommodate various energy management strategies, crucial for realistic WBAN implementations.

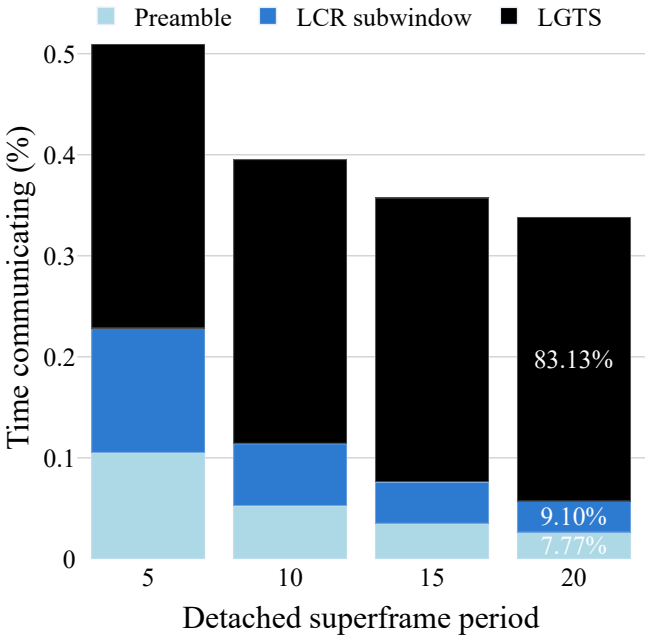

FIGURE 14. Average detached leaf time spent communicating, according to values in Table 4 and Table 5.

C. FULL CLUSTER SIMULATION

1) Materials and Methods

In this subsection, we quantitatively compare HB-MAC with IEEE 802.15.4 [5] in a well defined and realistic scenario. The performance evaluation is conducted by simulating both protocols with Castalia [29].

The simulation focuses on the intra-cluster operation, as this is HB-MAC's most innovative element, and this topology is simple enough to provide a fair performance comparison. Its aim is twofold. First, it demonstrates the performance of HB-MAC in terms of latency and energy consumption in comparison to the widely-used IEEE 802.15.4 standard. Second, it validates the effective operation of HB-MAC in a realistic setting including diverse data generation rates.

We simulated a single cluster composed of a hub surrounded by 6 leaves, supporting applications with diverse purposes summarized in Table 6.

The *packet generation rate* is an application layer parameter only, independent from the MAC layer.

A *communication period* is defined as the time interval between two leaf-hub data exchanges. An attached leaf selects its own communication period during the attachement procedure. However, detached leaves rely on the detached superframe period (DSP), a network-wide setting.

TABLE 6. Simulated leaves properties

Leaf	Data generation rate	Packet generation rate	Status	Communication period
1	24 bps	90 B every 30 s	Detached	DSP[a]
2	24 bps	3 B every 1 s	Detached	DSP[a]
3	120 bps	15 B every 1 s	Detached	DSP[a]
4	120 bps	15 B every 1 s	Attached	10
5	240 bps	6 B every 200 ms	Attached	10
6	240 bps	6 B every 200 ms	Attached	5

[a]Detached Superframe Period defined in Table 7

The chosen parameters for the HB-MAC protocol are listed in Table 7 (except otherwise stated).

TABLE 7. HB-MAC simulation parameters

Back-off strategy	UB
LGTS payload	6000 bits
Number of LCRTS	30
Number of DLGTS	3
Detached Superframe Period (DSP)	10
Heart Rate	[40, 80, 120, 160]

The chosen parameters for the IEEE 802.15.4 standard are listed in Table 8:

TABLE 8. IEEE 802.15.4 standard simulation parameters

Guard Time	1.5 ms
Max Payload	960 bits
Base Slot Duration	82 bits
Number of Superframe Slots	16
Beacon Order (BO)	6
Superframe Order (SO)	4
Unit Backoff Period	20 bits
Backoff Exponent (BE)	min 5 / max 7
Max CSMA Backoffs	4
Max Frame Retries	2
GTS length	2 slots per attached leaf

The physical layer configuration is the same for both protocols. The associated power consumption are listed in Table 9.

TABLE 9. Hardware power consumption

Parameters	Power Consumpion
Transmitter (Tx)	50 µW (for -40dBm) [28]
Receiver (Rx)	100 µW
Sleep Mode	1 µW [30]
Heartbeat detector (HB-MAC only)	58 nW [31]

The simulation starts with all the leaves in a reset state: they have no data to transmit and are unaware of their surroundings. It runs for a duration of 6000 s.

2) Simulation Results

Figure 15 shows the energy efficiency per useful bit (i.e. payload bit) delivered, for each leaf node in the network. Both the IEEE 802.15.4 standard and HB-MAC are simulated with a heart rate varying between 40 and 160 bpm.

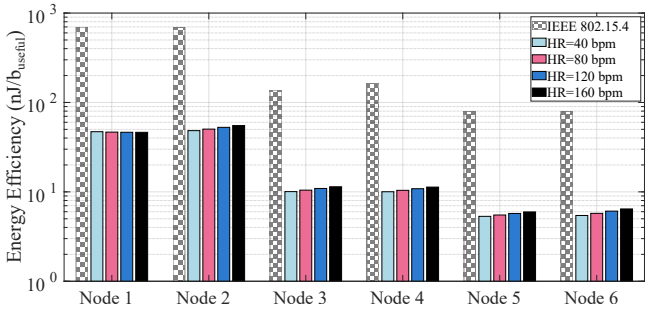

FIGURE 15. Energy efficiency per useful bit (payload) for both communication protocols and multiple nodes settings.

Overall, compared to the IEEE 802.15.4 standard, HB-MAC lowers the per-bit consumed energy by a factor between 12x and 16x depending on the considered leaf node. In addition, both protocols become more energy-efficient as the sensor data generation rate increases. This is explained by the minor contribution of packet transmission to the total energy consumption. The latter is dominated by idle listening or sleeping. Consequently, increasing the data generation rate improves the energy efficiency. Heart rate variation has a negligible impact on the energy efficiency.

Comparing performance in pairs of nodes with identical data generation rates allows to estimate the influence of the other communication parameters. Observing nodes 1 and 2, the emission of more frequent but smaller data packets improves the energy efficiency of IEEE 802.15.4, and degrades that of HB-MAC: smaller packets imply more communication overhead. The radio also consumes less in transmit than in receive, which is the default mode for IEEE 802.15.4. Finally, transmitting more decreases the energy consumption. For HB-MAC, the radio is in sleep mode by default, so more energy will be consumed with smaller packets.

Contrasting nodes 3 and 4 tells us that attached operation incurs a higher cost to IEEE 802.15.4, but slightly improves energy efficiency of HB-MAC. The overhead of the former comes from idle listening in both detached and attached mode during the whole random-access period, even though the attached leaf will communicate in its GTS. The slight improvement of HB-MAC is due to not having to reserve a slot before every guaranteed communication. Comparing nodes 5 and 6 shows that increasing the communication frequency of an attached leaf increases HB-MAC's consumption, but the additional cost is relatively low (< 8%).

Figure 16 presents HB-MAC's latency statistics as a function of the heart rate and the chosen detached superframe period. The latency of IEEE 802.15.4 is also given for comparison. This analysis is done for a hub surrounded by three identical detached leaves, generating 120 bps. The error bars correspond to the interval $[\bar{x} - \sigma, \bar{x} + \sigma]$, with $\bar{x}$ the average latency and σ the standard deviation.

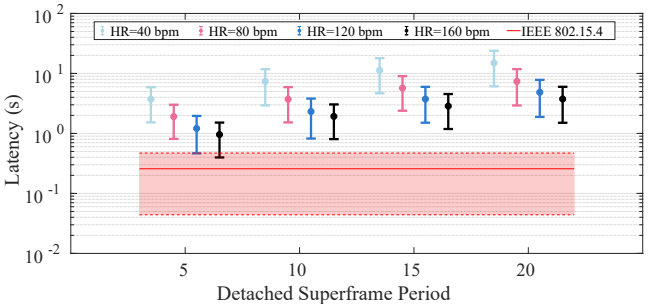

FIGURE 16. Communication latency for both communication protocols as a function of the detached superframe period.

The average latency with HB-MAC can range from 1 to 15 s, which is significantly longer than the sub-1s latency result obtained by 802.15.4. The obtained result also greatly varies with heart rate and communication frequency. As expected, the observed latency increases with the detached superframe period, since a longer period means a lower frequency of communication for the detached leaves. A higher heart rate implies shorter superframes and thus increased frequency of communication when considering a fixed detached superframe period: latency is decreased with physical activity of the user.

The simulations set up in this subsection clearly show that HB-MAC improves the leaf energy efficiency significantly when compared to a wide-spread standard such as IEEE 802.15.4. This comes at the price of a generally degraded latency, although tighter latency requirements can be met by operating in attached mode or using the emergency signaling mechanisms.

D. HEARTBEAT DETECTOR MANAGEMENT

Integrated circuit heartbeat detector state-of-the-art implementation reaches a power consumption as low as 58 nW [31]. Regardless of this ultra-low consumption, it remains non-negligible for a leaf given the targeted application. This section examines a duty-cycling scheme that aims to decrease the leaves' heartbeat detection energy cost. The envisioned

duty cycling strategy consists of turning ON the heartbeat detector a short period of time γ ahead of the next expected heartbeat as represented in Fig. 17. The expected inter-heartbeat duration is assumed identical to its previous occurrence.

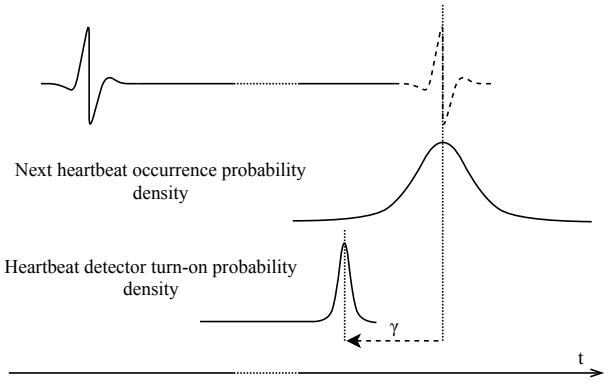

FIGURE 17. Heartbeat detector duty cycle strategy principle.

Two sources of uncertainty are evaluated: the leaf clock uncertainty and the heartbeat local variability. They are modeled following two random variables: X, representing the heartbeat detector start time; and Y, representing the actual heartbeat occurrence time. The time reference is set to the expected heartbeat time of occurrence.

X and Y are modeled with independent normal distributions $N(\mu_X = -\gamma, \sigma_X^2 = 4\mu s^2)$ and $N(\mu_Y = 0, \sigma_Y^2 = 900\mu s^2)$. The standard deviations are respectively chosen according to the maximal clock inaccuracy given in [7], and the heartbeat local variability described in [12]. Hence, the difference $Y - X$ follows a normal distribution centered in γ of variance $\sigma_X^2 + \sigma_Y^2$.

Optimizing γ is essential. Chosen too big, the detector will start too early and energy will be wasted. Chosen too small (including negative), the detector will miss the heartbeat. The heartbeat detection failure cost is modeled with the heartbeat detector staying ON for an entire superframe. In addition, the leaf will listen to the countdown slot following the next successful heartbeat detection. We calculate E_{cdwn} = 154 nJ/read the transceiver energy cost for a single countdown read operation, taking into account half of the corresponding guard interval.

The expected ON-time ratio $E[d_{on}]$ is detailed in (13). It assumes that the heartbeat detector can be switched ON and OFF without associated cost and that no heartbeat is missed when the heartbeat detector is ON.

$$E[d_{on}] = P_{hbd} \int_{-\gamma}^{+\infty} \frac{t}{T} \phi_\gamma(t) dt + \frac{1}{2} \Phi_\gamma(0) (P_{hbd} \int_{-T-\gamma}^{+\infty} (1 + \frac{t}{T}) \phi_{\gamma+T}(t) dt + E_{cdwn}) \quad (13)$$

For robustness purposes, a misdetection identification feature is required. If the current heartbeat duration exceeds an upper limit (set as multiple standard deviations σ_Y), the heartbeat is considered missed.

From a leaf standpoint, a missed heartbeat only has a limited impact, as long as it does not directly precede a superframe where communication was intended. A hub misdetection would however block the communication for the entire superframe. Therefore, hubs should not implement such an aggressive heartbeat detector duty cycling scheme and should prefer more robust and reliable schemes, relying on a lower bound of heartbeat interval durations.

Fig. 18 shows the leaf's average heartbeat detector power consumption with the duty cycling scheme activated. It achieves an energy saving from 74.7% at 210 bpm up to 94.7% at 36 bpm in comparison to an always-ON heartbeat detector.

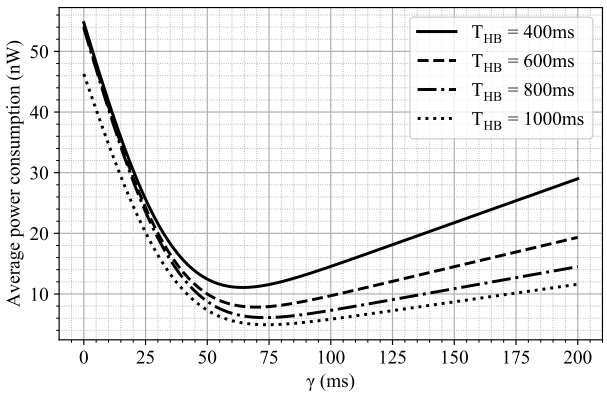

FIGURE 18. Average leaf duty-cycled heartbeat detector power consumption, including the possible additional countdown slot to read, with a reference power consumption of 58 nW.

V. DISCUSSION AND COMPARISON

A. RELEVANCE TO THE HUMAN BODY

Based on the heartbeat as a system-wide clock and on C-BCC as a physical layer, HB-MAC has been designed for energy-efficient communication on the human body, interfacing leaves and hubs.

Physical channel interruptions between hubs are mitigated by the use of pre-existing MAC layer standards, out of the scope of this work. Intra-cluster channel interruptions are mitigated by using a countdown slot and scheduling the intra-cluster communication right after it in the superframe structure. Countdown slots have a triple interest. They detect a channel deterioration and permit leaf re-synchronization in case of heartbeat misdetection or long channel interruptions. Ultimately, they allow leaves to notice a parent hub change, which is particularly important in attached mode.

Section IV-A shows the protocol's resilience to unusual traffic loads, which may follow long channel interruptions, treated with reasonable failure rate and energy cost. To further minimize the failure rate in random-access procedures after channel interruption, leaves could adapt their random-access strategy to the estimated number of concurrent leaves.

Those adaptation mechanisms are beyond the scope of this work.

The Attached operation is beneficial to upper-layer protocols for regular traffics, as it provides leaves with a specific address within the network. This is also applicable to detached leaves, but the address may be unstable. Downlink traffic addressed to leaves is a crucial feature for human body networks. It provides flexibility by allowing, for instance, sensor sampling rate and behavior update, providing better measurements and higher energy efficiency. On the other hand, efficient detached operation is a strong feature which efficiently supports event-driven and irregular traffic.

A WBAN supporting diverse applications inherently contains a wide range of devices and data rates as listed in Table 1. The HB-MAC protocol effectively supports the networking of devices with data generation rates ranging from some bits to multiple megabits per second. It integrates generic protocols into a framework tolerating detached leaf operation and supporting imbalanced device capabilities by design.

Finally, a WBAN must manage emergency events efficiently. Emergency mechanisms must have minimal impact on the regular network operation, but need a low response latency. HB-MAC implements an emergency signaling mechanism with a latency of not more than one heartbeat period and can detect heart failures.

B. COMPARISON

Multiple existing protocols address media access control in WBANs and WPANs. This section compares HB-MAC with the IEEE 802.15.4 [5], IEEE 802.15.6 [4] and Bluetooth Low Energy (BLE) standards [32], with H-MAC presented in [6], the low duty cycle protocol presented in [8] and the LDC-MAC protocol presented in [9]. The comparison is summarized in Table 10.

H-MAC is another MAC protocol based on the heartbeat as a universal clock, whereas [8] presents a protocol whose topology and device hierarchy relate to HB-MAC. LDC-MAC relates to HB-MAC as it supports cohabitation between devices featuring various data rates.

Most protocols targeting low-power operation rely on time division, often including random-access. Among the comparison references, only the protocols presented in [6] and [8] do not offer random-access, while LDC-MAC only offers a random-access period optionally, which is activated in case of a frame collision. The gain in terms of energy efficiency is effective when not offering random-access, however this design choice impairs flexibility. In particular, it does not feature efficient support for event-driven applications. Therefore, HB-MAC only uses random access in detached mode. In the present work, we characterized costs of random-access and sized it for different situations to offer flexibility at a reasonable cost. BLE, on the other hand, relies on frequency diversity, which is not offered by UWB-based physical layers, which attract a lot of attention [33].

TABLE 10. Protocols comparison

MAC protocol	HB-MAC [This work]	H-MAC [6]	Low Duty Cycle WBAN [8]	LDC-MAC [9]	IEEE 802.15.4 [5]	IEEE 802.15.6 [4]	BLE [32]
MAC access mechanism	Slotted CSMA/CA, with GTS	TDMA	TDMA	TDMA, optionally slotted CSMA/CA	CSMA/CA, optionally slotted, optional GTS	Slotted CSMA/CA, optional GTS	Frequency hopping
Topology	Mesh of stars	Star	Two-tier tree	Star, extensible to mesh	Star or peer-to-peer	Star	Star, tree or mesh
Synchronization	Yes - Biosignal-generated	Yes - Biosignal-generated	Yes	Yes	Optional	Optional	Broadcasting for discovery
Network scalability	High	Low	Medium	High	High	Medium	High
Adaptability to topology changes	Medium	High	Low (manual)	Medium	High	High	High
Emergency management strategy	Post-heartbeat slot	None	Alarm bit in data packet	Multiple DATA sections in insertion slots	None	Dedicated contention period for emergency	None
Connectionless operation possible	Yes	No	No	Yes	Yes	Yes	Partially
Traffic handling	Wide data rate span: kb/s-Mb/s, event-driven transactions	One device per heartbeat only, data rate: 250 kb/s[a]	Inflexible communication time distribution, data rate: 20-200 kb/s[b]	Wide data rate span, including emergency messages	Various traffic loads, data rate up to 250 kb/s	Various traffic loads, data rate up to multiple Mb/s	Various traffic loads, data rate up to 1 Mb/s
Energy consumption overhead	random-access for reservation & countdown slots	Scheduling & (re)synchronization packets	Beacon generation	Beacon generation	Optional beacon generation random-access for data	Optional beacon generation random-access for data	Advertising

[a] In the experimental setup, but the coordinator can adapt the scheduling to data rate requirements
[b] Several setups are compared in terms of power consumption

Mesh topologies offer robustness and become necessary for propagating a message across the whole body when relying on lossy physical layers, unable to reach the human body endpoints in one or two hops [10]. The protocol presented in [8] takes advantage of the limited range to isolate intra-cluster communications. Despite the similarities with this work, [8] limits its topology to a star of master nodes (equivalent to HB-MAC's hubs). Each of those masters is simultaneously the center of a local star network, made of sensor nodes (equivalent to HB-MAC's leaves). HB-MAC extends this topology to a mesh of hubs, each surrounded by a star of leaves, although the mesh management is delegated to an upper network layer.

Beacon synchronization is usually used to lower the sensor nodes' energy consumption. Several protocols, such as the IEEE 802.15.4 and IEEE 802.15.6 standards, make these optional. Similarly to H-MAC, the HB-MAC protocol uses heartbeat signals as out-of-band beacons. In a heartbeat-synchronized network, devices do not need to listen to an in-band beacon, potentially bringing down the communication system idle listening. This heartbeat-based beacon outperforms the LDC-MAC trade-off between coordinator power consumption, which decreases with the superframe length, and communication latency.

As a major requirement of modern WBAN applications, HB-MAC supports a wide span of data rates, up to multiple Mbps between hubs, depending on the inter-hub protocol implementation and the trade-offs made between hub-to-hub dedicated time and leaf-dedicated time. Such high data rates are not typical for low-power WBANs, therefore the only comparison references featuring similar maximal data rates are the IEEE 802.15.6 standard [4] and to some extent [8] and [9]. HB-MAC is consequently compatible with more demanding applications such as neural recording or smart prosthetic control. It is capable of flawlessly supporting high peak demands.

Body dynamics can modify the network geometry. While strong path loss permits cluster isolation in reference [8] and HB-MAC, it poses a specific challenge. In the case of C-BCC, for instance, movements can substantially alter the network topology [23]. The former shows poor resilience to topology changes, as the TDMA schedule is fixed. When it comes to HB-MAC, the resilience of the inter-hub mesh has not been precisely addressed as it relies on the inter-hub MAC protocol. The intra-cluster protocol has however been designed to be resilient, as detached superframes are globally scheduled and sized. In addition detached access is the same in all clusters. The resilience of the attached operation is granted by the hub address field in the countdown message. A leaf detecting a change in the hub address goes into detached mode to initiate a new attachment procedure. It avoids interference with the new cluster. As opposed to [8] and HB-MAC, protocols with a simpler topology naturally endure device movements with more ease.

VI. CONCLUSION

In this paper, we presented HB-MAC, a MAC protocol dedicated to WBANs relying on a heartbeat-based synchronization and compatible with Body Coupled Communication. It addresses important challenges of body area networks by supporting devices with imbalanced capabilities, a wide span of data rates, including regular as well as event-driven traffic.

We first presented a candidate physical layer, C-BCC on ultra-wide-band, adapted to low-power operation on the human body. Then, we described the MAC protocol architecture. As many low-power protocols, HB-MAC relies on a time division between random-access and guaranteed time slots. It has been designed for efficient operation under a heartbeat-based universal clock and supports an efficient integration of lightweight devices with irregular traffic.

We quantitatively addressed the energy cost of random-access for the case of detached operation including attachment procedures and collision probability during random-access. We deduced the optimal random-access strategy depending on the number of available time slots and concurrent leaves denoting the total number of leaves in a cluster along with their traffic. We deduced that the integration of efficient random-access algorithms allows HB-MAC to efficiently support event-driven and irregular traffic.

We also evaluated HB-MAC's performance in comparison with the IEEE 802.15.4 standard, by performing a network simulation of a realistic Human Intranet scenario. The performance analysis shows that, at the cost of increased latency, HB-MAC outperforms IEEE 802.15.4 by a factor between 12x and 16x in terms of energy efficiency.

To keep reducing the power consumption, we introduced a heartbeat detector duty-cycling scheme for leaves, proper to HB-MAC, and estimated the energy cost of heartbeat detection for leaves.

As this paper focused on the most innovative part of HB-MAC, which is the integration of heartbeat as a universal clock and joint support for regular and event-driven traffic, the integration of precise protocols for attached leaf communication and inter-hub communication remains as future work. The development of the upper layers of the network, as well as efficient distributed algorithms providing additional flexibility in the HB-MAC protocol with low overhead, remain as future work as well.

REFERENCES

[1] J. M. Rabaey, "The Human Intranet–Where Swarms and Humans Meet," IEEE Pervasive Computing, vol. 14, no. 1, pp. 78–83, 2015.

[2] A. Rahim, N. Javaid, M. Aslam, Z. Rahman, U. Qasim, and Z. A. Khan, "A Comprehensive Survey of MAC Protocols for Wireless Body Area Networks," in 2012 Seventh International Conference on Broadband, Wireless Computing, Communication and Applications. IEEE, 2012, pp. 434–439.

[3] F. Ullah, A. H. Abdullah, O. Kaiwartya, S. Kumar, and M. M. Arshad, "Medium access control (mac) for wireless body area network (wban): Superframe structure, multiple access technique, taxonomy, and challenges," Human-centric Computing and Information Sciences, vol. 7, no. 1, p. 34, 2017.

[4] IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Part 15.6: Wireless Body Area Networks," IEEE Std 802.15.6-2012, pp. 1–271, 2012.

[5] IEEE, "IEEE Standard for Low-Rate Wireless Networks," IEEE Std 802.15. 4-2015 (Revision of IEEE Std 802.15. 4-2011), 2016.

[6] H. Li and J. Tan, "Heartbeat-driven Medium-Access Control for Body Sensor Networks," IEEE transactions on information technology in biomedicine, vol. 14, no. 1, pp. 44–51, 2009.

[7] R. Benarrouch, A. Moin, F. Solt, A. Frappé, A. Cathelin, A. Kaiser, and J. Rabaey, "Heartbeat-Based Synchronization Scheme for the Human Intranet: Modeling and Analysis," arXiv preprint arXiv:2005.05915, 2020.

[8] S. J. Marinkovic, E. M. Popovici, C. Spagnol, S. Faul, and W. P. Marnane, "Energy-efficient Low Duty Cycle MAC Protocol for Wireless Body Area Networks," IEEE Transactions on Information Technology in Biomedicine, vol. 13, no. 6, pp. 915–925, 2009.

[9] C. Zhang, Y. Wang, Y. Liang, M. Shu, J. Zhang, and L. Ni, "Low duty-cycling mac protocol for low data-rate medical wireless body area networks," Sensors, vol. 17, no. 5, p. 1134, 2017.

[10] R. Benarrouch, A. Thielens, A. Cathelin, A. Frappé, A. Kaiser, and J. Rabaey, "Capacitive Body-Coupled Communication in the 400–500 MHz Frequency Band," in EAI International Conference on Body Area Networks. Springer, 2019, pp. 218–235.

[11] N. Choudhury, R. Matam, M. Mukherjee, and L. Shu, "Distributed beacon synchronization mechanism for 802.15. 4 cluster-tree topology," in International Wireless Internet Conference. Springer, 2016, pp. 10–20.

[12] U. R. Acharya, K. P. Joseph, N. Kannathal, C. M. Lim, and J. S. Suri, "Heart Rate Variability: a Review," Medical and biological engineering and computing, vol. 44, no. 12, pp. 1031–1051, 2006.

[13] H. Tanaka, K. D. Monahan, and D. R. Seals, "Age-predicted Maximal Heart Rate Revisited," Journal of the american college of cardiology, vol. 37, no. 1, pp. 153–156, 2001.

[14] P. Bjerregaard, "Mean 24 hour Heart Hate, Minimal Heart Rate and Pauses in Healthy Subjects 40–79 Years of Age," European Heart Journal, vol. 4, no. 1, pp. 44–51, 1983.

[15] T. Buchner and J. Gierałtowski, "How Fast Does the ECG Signal Propagate within the Body," Working Group for Cardiovascular Physics. Faculty of Physics, Warsaw University of Technology. Sixth Cardiology Meets Physics & Mathematics, At Zakopane, vol. 6, 2015.

[16] J. A. Afonso, P. Macedo, L. A. Rocha, and J. H. Correia, "Hierarchical Wireless Networks of Body Sensor Networks for Healthcare Applications," in Handbook of Research on Developments in E-Health and Telemedicine: Technological and Social Perspectives. IGI Global, 2010, pp. 65–86.

[17] Y. Xue and L. Jin, "A Naturalistic 3D Acceleration-based Activity Dataset & Benchmark Evaluations," in 2010 IEEE International Conference on Systems, Man and Cybernetics. IEEE, 2010, pp. 4081–4085.

[18] A. Ahmadi, D. D. Rowlands, and D. A. James, "Investigating the Translational and Rotational Motion of the Swing using Accelerometers for Athlete Skill Assessment," in SENSORS, 2006 IEEE. IEEE, 2006, pp. 980–983.

[19] A. Thielens, R. Benarrouch, S. Wielandt, M. G. Anderson, A. Moin, A. Cathelin, and J. M. Rabaey, "A Comparative Study of On-Body Radio-frequency Links in the 420 MHz–2.4 GHz Range," Sensors, vol. 18, no. 12, p. 4165, 2018.

[20] D. Marchaland, M. Villegas, G. Baudoin, C. Tinella, and D. Belot, "Novel Pulse Generator Architecture Dedicated to Low Data Rate UWB Systems," in The European Conference on Wireless Technology, 2005. IEEE, 2005, pp. 229–232.

[21] Q. Gu, RF System Design of Transceivers for Wireless Communications. Springer Science & Business Media, 2006.

[22] IEEE Standards Association and others, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE std, vol. 802, p. 2012, 2012.

[23] A. Moin, A. Thielens, A. Araujo, and J. M. Rabaey, "Adaptivity to enable an efficient and robust human intranet," arXiv preprint arXiv:1807.09723, 2018.

[24] R. Zheng, J. Hou, and L. Sha, "Performance Analysis of the IEEE 802.11 Power Saving Mode," in Proc. CNDS, 2004.

[25] N. Sornin, M. Luis, T. Eirich, T. Kramp, and O. Hersent, "LoraWAN Specification," LoRa alliance, 2015.

[26] N. Seyed Mazloum, "Body-coupled communications-experimental characterization, channel modeling and physical layer design," Master's thesis, Chalmers University of Technology, 2008.

[27] N. Cho, J. Yoo, S.-J. Song, J. Lee, S. Jeon, and H.-J. Yoo, "The Human Hody Hharacteristics as a Signal Transmission Medium for Intrabody Communication," IEEE transactions on microwave theory and techniques, vol. 55, no. 5, pp. 1080–1086, 2007.

[28] H. Cho, H. Kim, M. Kim, J. Jang, Y. Lee, K. J. Lee, J. Bae, and H.-J. Yoo, "A 79 pJ/b 80 Mb/s full-duplex transceiver and a 42.5 μW 100 kb/s super-regenerative transceiver for body channel communication," IEEE journal of solid-state circuits, vol. 51, no. 1, pp. 310–317, 2015.

[29] A. Boulis, "Castalia: revealing pitfalls in designing distributed algorithms in wsn," in Proceedings of the 5th international conference on Embedded networked sensor systems, 2007, pp. 407–408.

[30] X. Wang, J. Van den Heuvel, G.-J. van Schaik, C. Lu, Y. He, A. Ba, B. Busze, M. Ding, Y.-H. Liu, N. Winkel et al., "A 0.9–1.2 v supplied, 2.4 ghz bluetooth low energy 4.0/4.2 and 802.15. 4 transceiver soc optimized for battery life," in ESSCIRC Conference 2016: 42nd European Solid-State Circuits Conference. IEEE, 2016, pp. 125–128.

[31] D. Da He and C. G. Sodini, "A 58 nW ECG ASIC with Motion-tolerant Heartbeat timing Extraction for Wearable Cardiovascular Monitoring," IEEE transactions on biomedical circuits and systems, vol. 9, no. 3, pp. 370–376, 2014.

[32] Bluetooth SIG, "Core specification v5. 1," 2019.

[33] M. Hämäläinen and J. Iinatti, Wireless UWB body area networks: Using the IEEE802. 15.4-2011. Academic Press, 2014.

FLAVIEN SOLT received the B.S.Eng and M.S. degree in computer science from École Polytechnique, Palaiseau, France in 2017 and 2019, respectively.

He is currently pursuing the M.S. in electrical engineering and information technology at ETH Zürich, Switzerland. As part of his M.S. in computer science, he has been a Research Intern at STMicroelectronics, Crolles, France and a Visiting Student Researcher at the Berkeley Wireless Research Center (BWRC), University of California Berkeley, Berkeley, USA in 2019. His research interests include low-power network protocols and embedded systems.

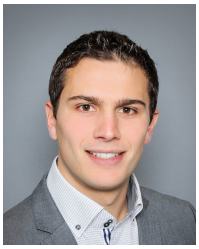

ROBIN BENARROUCH (S'17) received the B.S and M.S.Eng degree in Electrical Engineering from CPE Lyon, France in 2013 and 2015, respectively. He also received M.Res in Electronics and Embedded System from INSA Lyon - Centrale Lyon - UCBL, France in 2015.

He is currently pursuing his PhD with STMicroelectronics and the University of Lille 1, France, under a CIFRE program. He conducted half of his research as a Visiting Scholar Researcher at Berkeley Wireless Research Center (BWRC), University of California Berkeley, Berkeley, USA. His research interests include Body Coupled Communication and Low Power Wireless Communications for IoT and BAN applications.

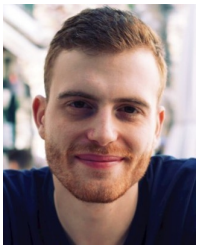

GUILLAUME TOCHOU (S'19) received his B.S and M.S. degree in electrical engineering from Grenoble INP - Phelma, France in 2016 and 2018, respectively.

He is currently pursuing his PhD through a CIFRE program at STMicroelectronics and University of Lille 1, France. As part of his PhD, he has been a Visiting Student Researcher at the Berkeley Wireless Research Center (BWRC), University of California, Berkeley. His research interests focus on ultra-low power wireless communication for IoT and Body Area Network applications.

OLIVÉR FACKLAM received the B.S.Eng degree in 2019, and is working toward the M.S. in computer science at École Polytechnique, Palaiseau, France.

In 2019, he has been an R&D Intern at Stereolabs. He is currently a Research Intern at STMicroelectronics, Crolles, France. His research interests include network protocols, embedded and autonomous systems.

ANTOINE FRAPPÉ (M'08-SM'17) graduated from the Institut Supérieur d'Electronique du Nord (ISEN), Lille, France in 2004. He received the M.Sc., Ph.D. and HDR (habilitation à diriger des recherches; French highest academic degree) in electrical engineering from the University of Lille, France in 2004, 2007 and 2019 respectively.

In 2007, he was with the Silicon Microelectronics group at the Institute of Electronics, Microelectronics, and Nanotechnologies (IEMN) in Villeneuve d'Ascq, France. He is now Associate Professor at Yncréa Hauts-de-France – ISEN Lille, in charge of the Electronics Team. His research interests concern digital RF transmitters, high-speed converters, mixed-signal design for RF and mmW communication systems, energy-efficient integrated systems, event-driven and neuro-inspired circuits for embedded machine learning.

He obtained a Fulbright grant in 2008 to pursue research in communication systems at the Berkeley Wireless Research Center (BWRC) at UC Berkeley, CA, USA. He was a co-recipient of the Best Student Paper Award at the 2011 Symposium on VLSI Circuits.

ANDREIA CATHELIN (M'04, SM'11) started electrical engineering at the Polytechnic Institute of Bucarest, Romania and graduated from ISEN Lille, France in 1994. In 1998 and 2013 respectively, she received PhD and "habilitation à diriger des recherches" (French highest academic degree) from the Université de Lille 1, France.

Since 1998, she has been with STMicroelectronics (STM), Crolles, France, now Fellow in Digital Front-End Manufacturing & Technology. Her focus areas are in the design of advanced RF/mmW/THz and ultra-low-power circuits and systems. Andreia has had numerous responsibilities inside IEEE since more than 15 years: at ISSCC, VLSI Symposium and ESSCIRC for TPC and Executive/Steering Committees respectively, and has been SSCS elected Adcom member 2015 to 2017.

Andreia is a co-recipient of the ISSCC 2012 Jan Van Vessem Award and of the ISSCC 2013 Jack Kilby Award and the winner of the 2012 STM Technology Council Innovation Prize, for having introduced on the company's roadmap the integrated CMOS THz technology for imaging applications. Very recently, Andreia has been awarded an Honorary Doctorate from the University of Lund, Sweden, promotion of 2020.

ANDREAS KAISER (S'85–M'87–SM'13) received the engineering diploma from the Institut Supérieur d'Electronique du Nord (ISEN), Lille, France, in 1984, and the PhD degree in 1990 from the University of Lille.

In 1990 he joined the Centre National de la Recherche Scientifique (CNRS) where he was responsible for the analog/RF IC design group at the Institut d'Electronique, de Microelectronique et de Nanotechnologies (IEMN) in Lille. He is currently vice-president for research and innovation at the Yncrea engineering graduate school. His research interests are continuous and discrete time analog circuits, data-converters, analogue design automation, RF-MEMS and RF circuits.

Prof. Kaiser served as TPC Chair of the European Solid State Circuits Conference in 1995 and 2005 and has been a guest and associate editor to the IEEE Journal of Solid State Circuits.

JAN M. RABAEY (S'80–M'83–SM'92–F'95) was a Research Manager with IMEC, Leuven, Belgium, from 1985 to 1987.

He is currently the Founding Director of the Berkeley Wireless Research Center (BWRC) and the Berkeley Ubiquitous SwarmLab. In 2019, he holds the CTO of the System-Technology Co-Optimization (STCO), Division of IMEC. He is also a Professor of the Graduate School, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley CA, USA, after being the holder of the Donald O. Pederson Distinguished Professorship for over 30 years. He has made high-impact contributions to a number of fields, including advanced wireless systems, low-power integrated circuits, mobile devices, sensor networks, and ubiquitous computing. His current research interests include the conception of the next-generation distributed systems, as well as the exploration of the interaction between the cyber and the biological world.

Prof. Rabaey is a member of the Royal Flemish Academy of Sciences and Arts of Belgium. He has received honorary doctorates from Lund (Sweden), Antwerp (Belgium), and Tampere (Finland). He has been involved in a broad variety of startup ventures, including Cortera Neurotechnologies, of which he was a co-founder. He was a recipient of major awards, among which the IEEE Mac Van Valkenburg Award, the European Design Automation Association (EDAA) Lifetime Achievement Award, the Semiconductor Industry Association (SIA) University Researcher Award, and the SRC Aristotle Award. He served as the Division Chair for the Electrical Engineering at Berkeley.