This is the author's version of an article that has been published in this journal.

Changes were made to this version by the publisher prior to publication.

The final version of record is available at https://doi.org/10.1109/WFCS46889.2021.9483595

# Enabling Listening Suspension in the Time Slotted Channel Hopping Protocol

Gianluca Cena, Stefano Scanzio, and Adriano Valenzano

National Research Council of Italy (CNR–IEIIT), Corso Duca degli Abruzzi 24, I-10129 Torino, Italy Email: {gianluca.cena, stefano.scanzio, adriano.valenzano}@ieiit.cnr.it

*Abstract*—Time slotted channel hopping provides reliable and deterministic communication in IEEE 802.15.4 mesh networks. Although slotted access is able to lower energy consumption drastically by reducing the duty cycle of the radio module, it usually leads to significant idle listening experienced by receivers, which makes it a sub-optimal solution when ultra low-power wireless is sought for.

In this paper a listening suspension mechanism is described, which operates at the MAC layer and is part of a more general approach aimed at cutting down energy consumption by proactively reducing idle listening. Links can be temporarily disabled, that convey slow-rate data streams whose characteristics, e.g., the generation period, are either known in advance to some extent or can be inferred by traffic inspection.

# I. INTRODUCTION

Time slotted channel hopping (TSCH) [1], [2] is an enhanced access mechanism for IEEE 802.15.4 [3] that, thanks to scheduled access to the transmission medium, offers high reliability, determinism, and, due to the reduced duty cycles, low power consumption [4], [5]. This enables the connection of battery-powered devices [6] in application scenarios like process and factory automation, e.g., to retrofit industrial plants in order to include new features and functions [7]. The medium access control (MAC) mechanism of TSCH relies on a periodically recurring slotframe, to which all nodes must be precisely synchronized. However, an asynchronous transmission service is offered to the users of the data-link layer, which is compatible with the Internet of Things (IoT) and Industrial IoT (IIoT) paradigms, where the Internet Protocol (IP) is exploited at the network layer. For this reason, some interesting IoT solutions have appeared recently, such as 6TiSCH [8], [9], [10], which rely on TSCH for data transmission.

In most wireless sensor networks (WSN) sensing is performed periodically by motes. Packet transmission in TSCH occurs cyclically too, but the sampling periods selected by applications are generally uncorrelated to the slotframe duration. Therefore, network configuration and setup of applications can be decoupled, and the design and deployment of distributed functions in heterogeneous systems (e.g., remote diagnostics and proactive maintenance) made easier. This means that two distinct kinds of periodicity can be found in TSCH-based solutions like 6TiSCH. At the MAC level, the duration  $T_{\rm sf}$ of the slotframe (in terms of number of slots) is selected network-wide. At the application level, instead, a number of

978-1-6654-2478-3/21/\$31.00 ©2021 IEEE

different sampling periods, denoted  $T_{c,i}$ , may be defined, each one depending on specific physical dynamics. For example, the temperature of a liquid in a tank may vary more slowly than a flow through a pipe. Clearly,  $T_{sf}$  must be chosen so that it is smaller than the shortest period for variables in the network,  $T_{sf} \leq \min_i \{T_{c,i}\}$ .

This is not the only constraint  $T_{\rm sf}$  must satisfy. Assuming that a single transmission opportunity (i.e., cell) is reserved in every slotframe for each pair of directly communicating nodes, to prevent network congestion the overall traffic on any link must never exceed the slotframe repetition rate, not even when interference and disturbance cause frame retransmissions. Network parameters should be set so as to take into account the probability that a transmission attempt fails. In practice, a safety margin has to be provided to avoid unwanted queuing phenomena in motes. In the current version of OpenWSN for OpenMote B, the retry limit *R* is set to 15, the default slot duration  $T_{\rm slot}$  to 20 ms, and the slotframe size to 101 slots (i.e.,  $T_{\rm sf} = 2.02$  s). Thus, sampling rates as low as half a minute can be quietly selected, network- and link-capacity wise.

In many applications, like condition-based monitoring involved in predictive maintenance, sampling periods can be much longer than the slotframe. If so, a non-negligible amount of energy is wasted in TSCH, which can shorten the operating time significantly when motes are powered on batteries. TSCH does not introduce any waste of energy on the transmitting side of a link. In fact, if there is no packet queued for transmission when a slot becomes available, no attempt is performed and the cell is simply left unused. By contrast, the receiving side of the link must be enabled in every cell it is associated, since it must be ready to get each frame potentially transmitted. As a consequence the receiver mote must be up and listening for some portion of the slot, in order to establish whether or not a transmission is being performed by some nearby peer. Unfortunately, in many real situations listening to the channel when a cell is unused is comparable to a frame reception from the power consumption viewpoint. This phenomenon is known as *idle listening*, and is the cause of energy wastes.

A simple solution, when all sampling periods are much larger than the slotframe duration, is the reconfiguration of the MAC protocol by increasing the number of slots in the slotframe ( $T_{\rm slot}$  depends on the physical layer and can hardly be changed). This introduces a drawback affecting cell  $\langle 0, 0 \rangle$ , that is the cell placed at slot and channel offsets 0 in the TSCH matrix, which is used by 6TiSCH for neighbor discovery

and RPL topology construction activities. In fact, network reconfiguration (e.g., when some motes are moved or switched on/off, obstacles are placed between them, nearby sources of interference and noise appear/disappear) becomes slower as  $T_{\rm sf}$  increases. In many cases this limitation is not severe since, from a practical viewpoint, having network reconfiguration times much shorter than the sampling period makes little sense. However, when sampling periods are very different (e.g., when they range from few minutes to several hours) enlarging  $T_{\rm sf}$  may not suffice. This is because  $T_{\rm sf}$  has to be selected starting with the highest dynamic packet flow, which makes this approach useful only in part for the slower streams. Moreover, excessive increases of  $T_{\rm sf}$  reduce the ability of nodes to keep their time sources synchronized, which is a basic requisite for correct time slotting operation.

A more effective solution to this problem, which enables ultra-low power communication, consists of suspending the receiving mote when no frame is expected to arrive, so that no energy is wasted because of idle listening. Such a kind of technique, named Proactive Reduction of Idle Listening (PRIL), was introduced in [11]. It is worth noting that acting this way is not trivial, as the involved motes must know in advance when transmissions occur, and this is in contrast with the IoT paradigm where packet transmission requests are driven by applications. A practical method is to explicitly drive the listening suspension (LS) of the link receiver by the transmitter, through the inclusion of suitable sleep commands in the frames being sent. This approach, which requires slight changes to the TSCH MAC that do not impair backward compatibility, was envisaged in [11] by exploiting suitable information elements, embedded in IEEE 802.15.4 frames, to put the receiver to sleep. In this way the transmitter holds all the knowledge about data exchanges, while the receiver simply obeys to sleep commands. Doing so prevents inconsistencies between link sides and tangibly improves LS reliability.

Mechanisms employed at the data-link layer to disable listening in selected cells should be clearly separated from the policies defined at higher levels to reduce the amount of idle listening. In this paper we focus on the former issue, by defining commands to be included in the MAC, as well as a number of relevant strategies. Our goal is to provide some hints and bounds on the benefits that LS can offer in realistic scenarios. Characterization of traffic patterns over links produced by applications and devices, and the analysis of optimal LS strategies are left for future work.

The paper is structured as follows: Section II introduces the problem of reducing idle listening, whereas in Section III sleep commands and some related strategies are described. Section IV focuses on power saving achievable with different strategies, while some numerical results are reported in Section V. Finally, Section VI concludes the paper.

# II. REDUCING IDLE LISTENING

Information about packet transfer timing can be obtained in two different ways: it is either provided *explicitly* by the user, who instructs the data-link layer about its traffic flow, or the data-link layer tries to *implicitly* infer the pattern of exchanges by continuously analyzing the traffic flowing through the mote. Both approaches are made harder as intermediate nodes in a mesh WSN forward packets produced by ascendants and descendants. In general, combining implicit and explicit data produces best results. Information about traffic can be obtained from different layers. For example, the network layer permits to identify streams coming from or directed to specific endpoints along a route. This is useful to decompose the traffic flowing through intermediate relays into separate and simpler contributions. In a similar way, application processes know the characteristics of the traffic they produce and consume (e.g., periodicity, minimum inter-arrival times, deadlines).

Effective implementations consist of two separate and cooperating elements: PRIL and Link Suspension Entity (LSE). PRIL concerns network management and mechanisms to establish dynamically at runtime when and how long a link can be safely suspended, in such a way that application constraints (e.g., latency, throughput) are not violated for the packet flow. PRIL can include:

- *Link traffic model* to describe traffic patterns conveyed through the link. For example, a single periodic stream is completely characterized by its period, while a superposition of periodic streams is identified by the set of related periods and phase offsets. A sporadic stream can be defined by means of its average and minimum interarrival times.
- *Interfaces* to protocol layers. The MAC layer interface is of utmost importance and enables interactions with the LSE. Another interface is offered to application processes, and is meant to support explicit LS configuration. For instance, it allows the selection of a specific traffic model (periodic, sporadic, etc.), the specification of its parameters (e.g., the sampling period of sensors), and the definition of additional constraints known by applications only (e.g., relative deadlines for sporadic messages). *Dissectors* are also envisaged for implicit traffic characterization. They enable analyses of information in protocol headers, which are encapsulated in frames sent over the link (e.g., IP source and destination addresses, UDP source and destination ports).



Repeated slotframe on air (synchronized time grid)

Fig. 1. Block diagram of PRIL techniques (interacting with the LSE).

 Algorithms (intelligence), possibly based on machine learning, to (semi-)autonomously detect traffic patterns from data mentioned above (i.e., finding a traffic model and its parameter values fitting the link) and optimally drive the LS mechanism in the MAC.

Instead, LSE is part of the MAC and consists of:

- LS-related *commands*, to be embedded in exchanged frames. They include requests issued by the transmitter to temporarily disable listening on the receiver and other elements involved in specific PRIL techniques.
- Protocol finite state machines and local state variables implementing the LS mechanism and supporting a set of pre-defined LS strategies.
- Data-link layer *primitives* to enable interactions between PRIL and LSE. They should be simple, flexible, and powerful enough to cope with different strategies.

# III. LISTENING SUSPENSION MECHANISM

The LS mechanism for TSCH belongs to the MAC layer. New primitives have to be defined in the MAC sublayer management entity (MLME) to interact with the LSE. For example, a generic MLME-LS-SET can be designed to enable/disable LS on a given link, to select a specific strategy, and to specify the strategy operating parameters. A detailed specification of new primitives cannot abstract from a stable definition of traffic models and PRIL interfaces, which are not available yet. In addition, the behavior of data transmission primitives of the MAC common part sublayer (MCPS), such as MCPS-DATA, must also be modified. On the one hand, in fact, procedures involved in listening suspension have to be executed when needed (e.g., triggered by either transmission requests or local timers and counters) while, on the other hand, traffic needs to be analyzed by PRIL starting from transmission requests (e.g., by setting function hooks to inspect frames on the link), to infer its characteristics.

In the current proposal sleep commands are defined as a part of the LSE and sent by the transmitter to the link receiver in order to force the latter to temporarily disable listening. They are conveyed in data frames by using information elements (IE), that are special fields that can be optionally included in IEEE 802.15.4 frames to carry ancillary information. IEs extend the basic protocol but preserve backward compatibility with existing devices and networks. In defining a new IE for sleep commands, denoted *sleep IE*, a streamlined encoding must be envisaged so as not to undermine energy saving these mechanisms are intended to achieve. In fact, the inclusion of IEs in frames increases the number of bits sent on air and, consequently, the power consumption on both the transmitting and receiving sides.

Every sleep command only concerns the link where the including frame was received. This is because any intermediate mote acting as a relay can have multiple links connected to its children and/or parent neighbors, and suspension must be managed independently for each one of them. It is worth noting that the LS mechanism works on both the link sides, by either putting them on hold or restoring the conventional TSCH behavior. This prevents the sender from performing transmission attempts when listening is temporarily disabled on the receiver. Of course, suitable variables have to be defined on the transmitting and receiving sides of every link to keep track of its LS state (they do not refer to the motes as a whole).

In the following subsections some possible options are briefly described for LS commands and for general-purpose LS strategies relying on them.

### A. Basic Sleep Command

The simplest way to instruct the receiver about the duration of the listening suspension is by specifying a *sleep time* in the sleep IE. This time is encoded as a small integer, denoted  $N_{\rm slp}$ , representing the number of slotframes where listening has to be suspended (note that the final decision about the suspension is left to the receiver, which ensures backward compatibility). Upon reception, the link is disabled during the following  $N_{\rm slp}$ slotframes, then it is enabled again, restoring the conventional TSCH operation. Fig. 2 shows that the time interval before the next usable cell of the link is  $(N_{\rm slp} + 1)T_{\rm sf}$  wide. In the figure, the number of cells that, in every slotframe, precede and follow the cell allocated to the link is fixed but arbitrary.

The special value  $N_{\rm slp} = 0$  informs the receiver that it must be listening in the next slotframe, and it is roughly equivalent to not including any sleep command. However, while the absence of sleep commands has clearly no effects, setting  $N_{\rm slp} = 0$  can be used to trigger some specific actions. For example, when multiple cells are reserved for the same link in the slotframe, this can enforce an intra-slotframe LS behavior by skipping all the remaining cells in the current slotframe. This feature, which is useful when overprovisioning is exploited for high-capacity links [12], is not considered in the remaining part of this paper. Values  $N_{\rm slp} \ge 1$  are instead meant to enforce inter-slotframe listening suspension.

Another point concerns the maximum extension of the sleep time interval. In our opinion 6 bits are enough to encode  $N_{\rm slp}$  for most real situations, leading to link suspension periods up to about two minutes. Allocating more bits to  $N_{\rm slp}$  does not bring noticeable advantages from the point of view of power consumption. In fact, the shared cell  $\langle 0, 0 \rangle$  remains always active in every slotframe, so that longer sleeping periods for other cells become irrelevant for energy saving (unless the mote is involved in several links). In addition, a receiving cell disabled



Fig. 2. Conventional TSCH vs. LS mechanism.

by a sleep command cannot be (easily) re-enabled before the re-activation time is reached. Responsiveness to alarms and operations taking place on demand, such as firmware upload or configuration over the air, worsens and may become unacceptable when too long sleep intervals are enforced.

LS works properly when appropriate values are selected for  $N_{\rm slp}$  at any time. For example, if the link bears a single periodic packet stream,  $N_{\rm slp}$  should be chosen to match the packet generation period. As mentioned before, characterization of the traffic model and its parameters is up to PRIL, since the MAC layer has to be kept as simple as possible. Suitable primitives can drive intelligence in collecting information: for instance, some parameters (e.g., deadlines) can be provided by upper layers while others (e.g., periods) can be inferred by analyzing the traffic over the link.

## B. Periodic LS Strategy

Let us first consider a simple LS strategy fitting the needs of periodic traffic completely defined by its period  $T_{\rm c}$ . Under the assumption that only one cell is reserved per link, network stability requires that  $T_{\rm c} > T_{\rm sf}$  (strict order applies, as transmission errors are unavoidable). Whenever a new frame transmission request is issued for the link (e.g., through the MCPS-DATA.request primitive), the transmitter starts a local *frame sleep* counter  $C_{\rm fr}$  for the frame. The counter is initialized to  $C_{
m fr}~=~\lfloor au_{
m c}
floor~\geq~1,$  where  $au_{
m c}~=~T_{
m c}/T_{
m sf}$  is the normalized transmission period. It is worth noting that rounding down allows to re-enable the link and, in particular, the receiver listening before a new packet becomes available on the transmitter, thus reducing both the access time and the average number of queued packets. Of course, if  $T_c$  is not a multiple of  $T_{\rm sf}$ , idle listening may occasionally take place when the receiver is re-enabled, as depicted in Fig. 3.

When the frame is sent, the value of  $N_{\rm slp}$  in the sleep IE is set to  $C_{\rm fr}$ , moreover the sleep command is included in the frame only if  $N_{\rm slp} > 0$ . Reception of a sleep command causes the receiver to set its own *reception sleep* counter  $C_{\rm rx}$  to  $N_{\rm slp}$ . Upon ack frame reception, the  $C_{\rm fr}$  value is also loaded into the *transmission sleep* counter  $C_{\rm tx}$ . At the very beginning of every slot allocated to the link  $C_{\rm fr}$  is decreased by one, provided



Fig. 3. Periodic LS strategy (by varying the actual period).

it is greater than 0 (counters contain non-negative values). The same is true, at the end of the slot, for  $C_{tx}$  and  $C_{rx}$ , assuming that they were not set in the same slot. Each link side is assumed to be either enabled when the corresponding counter ( $C_{tx}$  for the transmitter and  $C_{rx}$  for the receiver) is equal to 0, or disabled otherwise. Since slotframes are kept synchronized all over the network by TSCH,  $C_{tx}$  and  $C_{rx}$  are updated coherently on both sides of the link and their values are always the same. Should synchronization between motes be lost, all sleep counters are reset to 0, hence deactivating any ongoing listening suspension (and restoring conventional TSCH operation).

Frame delivery for a given transmission request may be occasionally delayed as shown, for example, to the fourth request in Fig. 2.b, because of a prior command with an excessively large sleeping period. This event is dealt with by the LS mechanism in this way: if, for any reasons, the frame is not sent at the first opportunity following the transmission request and it has to wait in the local queue,  $C_{\rm fr}$  is correctly decreased by the LS algorithm. Consequently, if the sleep command eventually reaches the receiver, the re-enable instant is not affected. This also copes with the relevant case of failed transmission attempts (detected by expiration of the ack timeout), where  $C_{\rm fr}$  is decreased by one on each retry. In this way, transmission errors affect the current packet in the same way as TSCH, whereas the next packet does not suffer from additional delays with respect to the case when no errors occur.

To help recovering from queuing phenomena, which may take place for instance when long bursts of errors are experienced, the sleep command is skipped if more than one frame concerning the same link is pending in the transmission queue, thus preventing the link capacity from being throttled in these conditions (which could lead to buffer overruns).

This strategy can also be used for quasi-periodic traffic, where packet generation takes place cyclically at a nominal rate but release times are subject to non-negligible jitters. If the time interval between two consecutive packet transmission requests exceeds  $T_c$ , receiver idle listening occurs. Conversely, when the interval is shorter than  $T_c$ , the early packet experiences some delay since the transmission can start only when the link is re-enabled (i.e.,  $C_{tx} = 0$  and  $C_{rx} = 0$ ).

#### C. Responsiveness

Responsiveness in TSCH is limited by its MAC mechanism, which is based on a continuously repeating slotframe [13]. As pointed out in [14], in absence of errors due to frame corruption on air, one-way transmission latency consists of two elements, namely access delay and transmission time. Access delay is the time the transmitter spends waiting for the slot assigned to the link, and is the largest contribution. Under the assumption that, at any time, at most one frame is queued for transmission for the link, it can be modeled as a random variable uniformly distributed between 0 and  $T_{sf}$ . The transmission time includes the offset at the beginning of the slot (*macTsTxOffset*) as well as the time taken for transmitting the data frame on air. Since it is always strictly shorter than  $T_{\rm slot}$  we will neglect it in the following. In this way the *worst-case transmission latency*  $T_{\rm wc}$  is approximately the same as the maximum access delay when no errors occur.

The access delay may grow sensibly when sleep commands are exploited. With respect to conventional TSCH the period between usable cells for a link (and access delays) increases from  $T_{\rm sf}$  to  $(N_{\rm slp} + 1)T_{\rm sf}$ , while energy spent by the receiver for listening to the channel decreases by about the same factor.

With strictly periodic patterns of packets and a proper selection of  $N_{\rm slp}$  (ensuring that the link is re-enabled just before a new packet becomes available for transmission) delays are roughly the same as TSCH. However, packets generated sporadically can experience access delays up to  $(N_{\rm slp}+1)T_{\rm sf}$ when they become ready just after the link suspension. To improve link responsiveness for sporadic packets,  $N_{\rm slp}$  can be set based on their *relative deadline*  $T_{\rm d}$ . In particular, to ensure that  $T_{\rm wc} \leq T_{\rm d}$ , again in the optimistic case of an ideal channel without errors,  $N_{\rm slp}$  has to be set so that  $N_{\rm slp} \leq \tau_{\rm d} - 1$ , where  $\tau_{\rm d} = T_{\rm d}/T_{\rm sf}$  is the normalized deadline. For example, with the network parameters we considered,  $N_{\mathrm{slp}}$  must not exceed 13 if latency must be shorter than 30 s. If the period is sensibly larger than the deadline  $(T_c \gg T_d)$ , the link is reenabled too early most times and the LS mechanism behaves less effectively.

## D. Slow Periodic LS Strategy

In slow periodic streams the period may exceed the maximum value allowed for the sleep command ( $64T_{\rm sf}$  for  $N_{\rm slp}$  6bit encoding). In such cases the LSE takes care of issuing sleep commands automatically to repeatedly restart the sleeping period in the receiver and make the suspension last the desired amount of time.  $N_{\rm slp}$  is equal to its maximum value (63) in all the sleep commands in the sequence except for the last one, which is set to  $(|\tau_c|-1) \mod 64$  and re-enables the link at the intended time. Overall, the sequence includes  $[\tau_c/64]$  frames. For example, if  $T_c$  is equal to 10 minutes (i.e.,  $N_{slp} = 296$ ), five frames are sent per packet, spaced by  $64T_{\rm sf}$ : the first one is the data frame, which includes the packet and a sleep command with  $N_{\rm slp} = 63$ , followed by three empty frames (with no payload) where  $N_{\rm slp}=63$ , and one final (empty) frame where  $N_{\rm slp} = 39$ . Should a new transmission request be issued before the above procedure is completed, the course of action is interrupted by the transmitter and the new data frame is sent as soon as the link is re-enabled.

The impact of this approach on power consumption is relatively low and is roughly equivalent to the energy needed for sending and receiving one frame every 64 slotframes (about two minutes). Most of these frames contain no actual data and only deliver sleep commands to the receiver. To further reduce power consumption a specific short format can be envisaged for *empty sleep frames* that only carry the sleep IE. A new MAC sleep command frame can be defined to this purpose but other solutions are possible as well. Note that from the energy viewpoint this contribution is rather negligible when compared to cell  $\langle 0, 0 \rangle$ , which is active in every slotframe, nonetheless it improves network responsiveness by imposing a reasonable upper bound  $T_{\rm wc} = 64T_{\rm sf}$  on transmission latency.

## E. Extended Sleep Command

Energy consumption can be reduced without impairing responsiveness too much, by means of an extended version of the sleep command we call *xsleep*. The relevant IE includes two parameters:  $N_{\rm slp}$  and  $N_{\rm snz}$ .  $N_{\rm slp}$  is the number of slotframes the link must be kept disabled. It specifies the *nominal* duration of the sleeping period and is managed the same way as in basic sleep commands. However, the range of allowed values for the sleep time can be increased noticeably in this case without latency is consequently worsened. For instance,  $N_{\rm slp}$ can be encoded with 12 bits and the sleeping period can last up to  $4096T_{\rm sf}$  (more than two hours). Selection of larger field sizes to enable longer sleeping periods is clearly possible.

The *snooze* time, denoted  $N_{\rm snz}$ , specifies how often the link must be temporarily re-enabled (*awakened*) during the nominal sleeping period, in order to flush packets sporadically generated in the meanwhile and left pending in the transmission buffer. In this case, on correct reception of the *xsleep* command an iterative procedure is started, where the link is disabled for  $N_{\rm snz}$  consecutive slotframes before being awakened for one slotframe. This behavior is automatically repeated by the LSE until the end of the nominal suspension period (as specified by  $N_{\rm slp}$ ) is reached. Clearly,  $N_{\rm snz}$  must be strictly lower than  $N_{\rm slp}$ .  $N_{\rm snz}$  encoded with 6 bits (e.g.,  $N_{\rm snz} \leq 63$ ) provides upper bounds for the access time similar to the basic sleep command. In practice, 3 bytes are enough to encode the *xsleep* command into the content of a sleep IE.

Identification of slotframes where the link must be awakened does not consider the instant when the sleep command reaches the receiver, but proceeds from the end of the sleeping period moving backwards. In other words, frames can be sent only when  $C_{\rm tx} \mod (N_{\rm snz} + 1) = 0$ . In this way delays in setting  $C_{\rm rx}$  caused, for instance, by transmission errors do not affect the wake-up synchronization on the two link sides. For example, if  $N_{\rm slp} = 58$  ( $T_{\rm c} \simeq 120 \, {\rm s}$ ) and  $N_{\rm snz} = 13$ ( $T_{\rm d} \simeq 30 \, {\rm s}$ ), the receiver temporarily wakes up at the 3-rd, 17-th, 31-st, and 45-th slotframes after the command is issued, and it is re-enabled at the 59-th slotframe.



Fig. 4. Extended periodic LS strategy (with and without queuing).

Fig. 4 assumes  $N_{\rm slp} = 9$  and  $N_{\rm snz} = 3$ , and shows that awakening occurs when counters equal 8 and 4. If needed, the *xsleep* command can be aborted at those time instants, for instance when queuing occurs as in Fig. 4b. If, for whatever reason, a number of sporadic packets become pending in the transmission buffer (e.g., as a consequence of prolonged interference on air), the transmitting LSE sends a new sleep command to reset the behavior of the receiver. A special *xsleep* command with  $N_{\rm slp} = 0$  and  $N_{\rm snz} = 0$  can be defined to this purpose, which reverts the behavior to conventional TSCH. Then, queued packets can be drained at the maximum link speed and LS operation resumed when the queue is empty.

LS strategies that rely on *xsleep* trade energy for responsiveness. In fact, the worst-case access delay (and  $T_{\rm wc}$  as well) is equal to the *snoozing period*  $(N_{\rm snz} + 1)T_{\rm sf}$ . Selecting a low value for  $N_{\rm snz}$  reduces  $T_{\rm wc}$ , but also increases power consumption on the receiver side consequently. Because of the large values allowed for  $N_{\rm slp}$  (latency now only depends on  $N_{\rm snz}$ ) the time between transmissions of *xsleep* commands can be as large as about two hours, which makes energy consumption completely negligible on the transmitter side.

# F. Extended Periodic LS Strategy

This strategy is suitable for a traffic pattern consisting of the superposition of a quasi-periodic stream with (average) period  $T_{\rm c}$  and sporadic streams with relative deadline  $T_{\rm d}$ . For simplicity we assume that sporadic packets, which should be served promptly, are issued in consequence of rare events. In this case, *xsleep* commands are conveyed inside data frames carrying periodic packets, and  $N_{\rm snz}$  and  $N_{\rm slp}$  are selected so that  $N_{\rm snz} = \lfloor \tau_{\rm d} \rfloor - 1$  and  $N_{\rm slp} = \lfloor \tau_{\rm c} \rfloor - 1$ .

Such a strategy closely resembles periodic LS, but the receiver wakes up regularly during the sleeping period to ensure bounded transmission latency to sporadic packets. Since *xsleep* enables large sleeping times, high efficiency can be reached. Moreover, should the inferred period be updated by PRIL, changes can be enforced by sending a new *xsleep* command as soon as the receiver awakens.

# **IV. PERFORMANCE ANALYSIS**

Energy consumed on the transmitting and receiving link sides can be evaluated for conventional TSCH and LS strategies. In the following we assume that no transmission error occurs, to simplify the analysis and focus on the intrinsic energy saving capability of the proposed solutions.

#### A. Power Consumption Model

Let  $E_{txd}$  and  $E_{rxd}$  be the amounts of energy spent for performing a single data frame transmission and reception attempt, respectively. In the following, we model these quantities by assuming that they depend on the frame size linearly. In particular,  $E_{txd} = E_{tx0} + e_{txB} \cdot L_{phy}$ , where  $L_{phy}$  is the overall number of bytes transmitted by the physical layer,  $e_{txB}$  is the energy required to send one byte, and  $E_{tx0}$  is a constant amount of energy additionally spent for every transmission attempt. For instance, for OpenMote B boards running Open-WSN and communicating through 6TiSCH,  $E_{tx0} = 7 \,\mu J$  and  $e_{txB} = 2 \,\mu J/B$  (see plots of Fig. 4 in [2]), thus the largest fraction of power directly depends on the number of bytes transmitted on air.

A very similar relation holds for reception. From the same set of experimental data we obtain  $E_{rx0} = 65 \,\mu\text{J}$  and  $e_{rxB} = 1.3 \,\mu\text{J}/\text{B}$ . Unsurprisingly, the energy per byte is lower than for transmission. However, a non negligible amount of energy is spent because listening must be enabled for some time preceding the instant when transmission is expected to begin.  $E_{rx0}$  is a mean value, and fluctuations can be experienced depending on the alignment precision of the transmitter and receiver time grids at any given time.

Energy needed to send and receive the ack frame (33 B) is  $E_{\rm txa} = 106 \,\mu {\rm J}$  and  $E_{\rm rxa} = 79 \,\mu {\rm J}$ , respectively. This also includes operations to finalize a frame exchange besides the transmission/reception on air. Finally, the energy spent every time idle listening occurs for the cell is  $E_{\rm lis} = 138 \,\mu {\rm J}$ .  $E_{\rm lis}$  is about twice  $E_{\rm rx0}$ , since the listening interval is centered around the expected start of frame transmission.

In absence of frame transmission and reception the mote drains about  $31.4 \,\mathrm{mW}$  (i.e.,  $628 \,\mu\mathrm{J}$  in every slot). This is due to the basic activity of the operating system, protocol stack and application processes, and concerns all board components besides the microcontroller. It is worth observing that such a high value is likely due to the fact that both the OpenWSN software and the OpenMote B hardware have not been optimized yet. Hopefully, this constant consumption will be lowered dramatically when industry-grade devices based on this technology are designed and implemented. For these reasons this contribution in not considered in the following analysis, which focuses only on energy spent for communication.

# B. Conventional TSCH

Let  $\Lambda_{\rm sf} = 1/T_{\rm sf}$  and  $\Lambda_{\rm c} = 1/T_{\rm c}$  be the slotframe repetition rate and the packet transmission rate, respectively. The transmitter power consumption for a packet stream sent over a conventional TSCH link is

$$P_{\rm t} = (E_{\rm txd} + E_{\rm rxa})\Lambda_{\rm c}.$$
 (1)

When only reception (and acknowledgment) of frames is considered, for the receiver we have

$$P_{\rm r0} = (E_{\rm rxd} + E_{\rm txa})\Lambda_{\rm c}.$$
 (2)

 $P_{\rm r0}$  equals the power spent by the receiver when an oraclebased LS strategy is adopted. Thanks to the oracle, listening is enabled only when the cell conveys a frame, without the need to obtain any information from the sender by means of sleep commands embedded in IEs. As a consequence,  $P_{\rm r0}$  provides a lower bound to power consumption for a given stream. The overall receiver consumption is given by

$$P_{\rm r} = P_{\rm r0} + E_{\rm lis} [\Lambda_{\rm sf} - \Lambda_{\rm c}], \qquad (3)$$

where the rightmost term refers to idle listening, which occurs when the cell allocated to the link is not used for data (hence, the difference between  $\Lambda_{sf}$  and  $\Lambda_{c}$ ).

In a TSCH/6TiSCH network operating correctly, additional frame exchanges are performed automatically in every slot-frame. This is the case of transmission/reception in the shared cell  $\langle 0, 0 \rangle$ , which is meant for network management activities. Despite their contribution to power consumption, they are not taken into account here because they are not correlated with data streams produced by applications and affected by the LS mechanism. Roughly speaking, power consumption for these frame exchanges can be expressed as  $P_0 = \bar{E}_{\langle 0,0 \rangle} \Lambda_{\rm sf}$ , where  $\bar{E}_{\langle 0,0 \rangle}$  is the average energy spent by the mote to carry out one (non-confirmed) transmission/reception attempt in cell  $\langle 0,0 \rangle$ . In practice, the value of  $\bar{E}_{\langle 0,0 \rangle}$  does not differ from  $E_{\rm txd}$ ,  $E_{\rm rxd}$ , and  $E_{\rm lis}$  significantly.

# C. Periodic LS Strategy

In periodic LS a basic sleep command is included in every data frame where  $N_{\rm slp} = \lfloor \tau_{\rm c} \rfloor - 1$ . For example, if  $T_{\rm c} = 30 \, {\rm s}$  then  $N_{\rm slp} = 13$ . In the following we will assume that  $T_{\rm c}$  can be inferred precisely by the intelligence of the PRIL module. Obviously, this represents the best case for the LS mechanism, and enables the computation of an upper bound for improvements this approach can offer.

Power consumption on the transmitting side increases slightly because of the sleep command

$$P_{\rm t}^{\rm P} = P_{\rm t} + L_{\rm slp} e_{\rm txB} \Lambda_{\rm c},\tag{4}$$

where  $L_{\rm slp}$  is the size in bytes of the encapsulating IE ( $L_{\rm slp} = 3$  B in our case). For typical frame size (90 B) the increase is negligible (~ 3%).

Data frames that carry sleep commands suspend listening for  $N_{\rm slp}$  slotframes when they are received correctly. Thus every frame prevents idle listening from occurring  $N_{\rm slp} = \lfloor \tau_{\rm c} \rfloor - 1$  times.

Hence

$$P_{\rm r}^{\rm P} = P_{\rm r0} + L_{\rm slp} e_{\rm rxB} \Lambda_{\rm c} + E_{\rm lis} \left[ \Lambda_{\rm sf} - \lfloor \tau_{\rm c} \rfloor \Lambda_{\rm c} \right] = P_{\rm r} + L_{\rm slp} e_{\rm rxB} \Lambda_{\rm c} - E_{\rm lis} \left[ \lfloor \tau_{\rm c} \rfloor - 1 \right] \Lambda_{\rm c}.$$
(5)

Again, the consumption increase due to sleep IEs is typically negligible. Instead, the contribution due to idle listening (rightmost term in the first line of the equation) is considerably lower than TSCH, because the multiplicative factor for  $E_{\rm lis}$  is always smaller than  $\Lambda_{\rm c}$ , and equal to 0 if  $T_{\rm c}$  is a multiple of  $T_{\rm sf}$ . In the latter case the difference in terms of power consumption between the oracle and the periodic LS strategy is constant and equal, overall, to  $L_{\rm slp}(e_{\rm txB} + e_{\rm rxB})\Lambda_{\rm c}$ , which, for  $L_{\rm slp} = 3$  B and  $T_{\rm c} = 30.3$  s, means  $0.33 \,\mu$ W only.

# D. Slow Periodic LS Strategy

For slow packet rates  $(T_c > 64T_{sf})$  empty sleep frames are automatically sent by the LSE on the link transmitter to enlarge the sleeping period without worsening responsiveness. Power consumption on both sides can be derived from the periodic case by including this additional contribution

$$P_{\rm t}^{\rm slP} = P_{\rm t}^{\rm P} + E_{\rm txe} n_{\rm emp} \Lambda_{\rm c}, \qquad (6)$$

$$P_{\rm r}^{\rm slP} = P_{\rm r}^{\rm P} + E_{\rm rxe} n_{\rm emp} \Lambda_{\rm c}, \tag{7}$$

where  $n_{\rm emp} = \lceil \tau_{\rm c}/64 \rceil - 1$  is the number of empty sleep frames inserted after every data frame. Because of our assumptions,  $n_{\rm emp}\Lambda_{\rm c} < \Lambda_{\rm sf}/64$ , which sets an upper bound to the additional consumption needed to keep LS alive.

### E. Extended Periodic LS Strategy

We have shown that when the link conveys both sporadic streams of packets with relative deadline  $T_{\rm d}$  and a periodic stream with period  $T_{\rm c}$  the best approach is to adopt *xsleep* commands where  $N_{\rm slp} = \lfloor \tau_{\rm c} \rfloor - 1$  and  $N_{\rm snz} = \lfloor \tau_{\rm d} \rfloor - 1$ . With this strategy, the receiver wakes up temporarily  $n_{\rm wup}$  times before listening is re-enabled, and  $n_{\rm wup} = \lceil (N_{\rm slp}+1)/(N_{\rm snz}+1) \rceil - 1 = \lceil \lfloor \tau_{\rm c} \rfloor / \lfloor \tau_{\rm d} \rfloor \rceil - 1$ . For example, if  $T_{\rm c} = 600 \, {\rm s}$  and  $T_{\rm d} = 30 \, {\rm s}$ , then  $N_{\rm slp} = 296$ ,  $N_{\rm snz} = 13$ , and the link is awakened 21 times during every sleeping period. For the sake of simplicity, we do not consider cases where  $T_{\rm c} > 4096T_{\rm sf}$ .

Power consumption on the transmitting side is almost the same as for the periodic strategy

$$P_{\rm t}^{\rm xP} = P_{\rm t} + L_{\rm xslp} e_{\rm txB} \Lambda_{\rm c},\tag{8}$$

however the size  $L_{xslp}$  of the sleep IE is slightly larger than the basic sleep command ( $L_{xslp} = 5 \text{ B}$ ).

On the receiving side

$$P_{\rm r}^{\rm xP} = P_{\rm r0} + L_{\rm xslp} e_{\rm rxB} \Lambda_{\rm c} + E_{\rm lis} \left[\Lambda_{\rm sf} - \left(\lfloor\tau_{\rm c}\rfloor - n_{\rm wup}\right)\Lambda_{\rm c}\right]$$
$$= P_{\rm r}^{\rm P} + (L_{\rm xslp} - L_{\rm slp}) e_{\rm rxB} \Lambda_{\rm c} + E_{\rm lis} n_{\rm wup} \Lambda_{\rm c}.$$
(9)

As expected, power consumption is higher than for basic sleep commands with the same value of  $N_{\rm slp}$ , because of the additional contribution of idle listening introduced by the awakening mechanism to ensure better responsiveness.

#### V. PERFORMANCE COMPARISON

To compare strategies we assume that both periodicity and relative deadlines are known (or correctly estimated by PRIL). The metric we consider is the power required to deal with a packet stream that transfers application data on a single link. Only contributions due to communication are taken into account, as other kinds of power drain by motes can be reduced considerably by exploiting careful hardware and software optimization. The frame size  $L_{\rm phy}$  is selected equal to 90 B, to resemble ICMP echo requests generated by ping commands in real networks. Three different packet transmission rates  $\Lambda_{\rm c}$  were considered, corresponding to periods equal to 30 s, 2 minutes and 10 minutes respectively, while  $E_{\rm txe} = 87 \,\mu$ J and  $E_{\rm rxe} = 117 \,\mu$ J for the slow periodic LS strategy (labeled "Basic (slow)" in Table I). These values were obtained by considering empty sleep frames encoded with 40 B.

Results are reported in Table I. TSCH shows noticeably higher consumption than the oracle on the receiving side. Tangible saving is possible with the periodic LS strategy

#### TABLE I

Power consumption on both sides of the link and worst-case transmission latency  $T_{\rm wc}$  for the different LS strategies by varying packet generation period  $T_{\rm c}$  and relative deadline  $T_{\rm d}$ .

$T_{ m c}$ / $T_{ m d}$	Strategy	$N_{\rm slp}$	N <sub>snz</sub>	$T_{\rm wc}$	$P_{\rm t}$	$P_{\rm r}$
(s) / (s)				(s)	(µW)	(µW)
30 / -	Oracle	-	-	2.02	8.8667	9.6000
30 / -	TSCH	-	-	2.02	8.8667	73.3168
30 / -	Basic	13	-	28.28	9.0667	13.6468
120 / -	Oracle	-	-	2.02	2.2167	2.4000
120 / -	TSCH	-	-	2.02	2.2167	69.5668
120 / -	Basic	58	-	119.18	2.2667	2.8993
120 / 10	eXtended	58	3	8.08	2.3000	19.0210
120 / 30	eXtended	58	13	28.28	2.3000	7.5210
600 / -	Oracle	-	-	2.02	0.4433	0.4800
600 / -	TSCH	-	-	2.02	0.4433	68.5668
600 / -	Basic (slow)	296	-	129.28	1.0333	1.2733
600 / 10	eXtended	296	3	8.08	0.4600	17.5177
600 / 30	eXtended	296	13	28.28	0.4600	5.3277
600 / 120	eXtended	296	58	119.18	0.4600	1.6477

based on sleep commands ("Basic") and improves with larger periods. Unfortunately, responsiveness is also impaired, and this could be a problem if sporadic packets were occasionally conveyed on the same link. The periodic LS strategy based on extended sleep commands ("eXtended") is a reasonable tradeoff and is suited to those cases where the period is much larger than the relative deadline. By means of two degrees of freedom (sleep and snooze times), it achieves interesting power saving yet keeping the MAC complexity at a minimum. In general, it retains the expected low-power consumption of TSCH on the transmitting side, whereas responsiveness and energy saving on the receiver side can be balanced on a per-link basis.

## VI. CONCLUSION

By layering time slotting and channel hopping atop IEEE 802.15.4, TSCH offers high reliability, determinism, and low power consumption, which makes it a good candidate when mesh WSNs compliant to the IIoT paradigm have to be deployed in industrial scenarios. Unfortunately, the access technique of TSCH (but the same also holds for similar protocols that rely on scheduled access) implies non-negligible power consumption by receivers due to idle listening, i.e., when the receiving circuitry is switched on but no one is sending frames on air.

Details about the PRIL proposal have been presented in this paper, in order to reduce idle listening phenomena by exploiting available information about data exchanges. In particular, we focused on the listening suspension (LS) mechanism, to be incorporated in the TSCH MAC, that enables the receiving side of a link to be put on sleep on demand. Besides basic strategies, suitable for periodic packet streams, extended solutions have also been introduced, which offer a reasonable trade-off between power consumption and responsiveness for sporadically generated packets thanks to the ability to repeatedly snooze the receiver. Formulas have been derived for power consumption estimation for both sides of a link in 6TiSCH, which were used to evaluate the behavior of a link in realistic operating conditions, by leveraging the energy model of a real device (OpenMote B running OpenWSN). Results confirm that the amount of saved energy is worth the adoption of LS. Future activities in this area will focus on the evaluation of LS mechanisms through simulation of realistic (and more complex) traffic patterns and the implementation of PRIL algorithms to automatically infer fitting traffic models.

#### REFERENCES

- D. De Guglielmo, B. Al Nahas, S. Duquennoy, T. Voigt, and G. Anastasi, "Analysis and Experimental Evaluation of IEEE 802.15.4e TSCH CSMA-CA Algorithm," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 2, pp. 1573–1588, 2017.
- [2] S. Scanzio, M. G. Vakili, G. Cena, C. G. Demartini, B. Montrucchio, A. Valenzano, and C. Zunino, "Wireless Sensor Networks and TSCH: A Compromise Between Reliability, Power Consumption, and Latency," *IEEE Access*, vol. 8, pp. 167 042–167 058, 2020.
- [3] IEEE, "IEEE Standard for Low-Rate Wireless Networks," *IEEE Std* 802.15.4-2020 (Revision of IEEE Std 802.15.4-2015), pp. 1–800, 2020.
- [4] F. Moreno-Cruz, V. Toral-López, A. Escobar-Molero, V. U. Ruíz, A. Rivadeneyra, and D. P. Morales, "treNch: Ultra-Low Power Wireless Communication Protocol for IoT and Energy Harvesting," *Sensors*, vol. 20, no. 21, 2020. [Online]. Available: https://www.mdpi.com/ 1424-8220/20/21/6156
- [5] V. Venkateswaran and I. O. Kennedy, "How to sleep, control and transfer data in an energy constrained wireless sensor network," in 2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2013, pp. 839–846.
- [6] H. Yetgin, K. T. K. Cheung, M. El-Hajjar, and L. H. Hanzo, "A Survey of Network Lifetime Maximization Techniques in Wireless Sensor Networks," *IEEE Communications Surveys Tutorials*, vol. 19, no. 2, pp. 828–854, 2017.
- [7] Z. Zhang, A. Mehmood, L. Shu, Z. Huo, Y. Zhang, and M. Mukherjee, "A Survey on Fault Diagnosis in Wireless Sensor Networks," *IEEE Access*, vol. 6, pp. 11349–11364, 2018.
- [8] P. Thubert, "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4," *IETF Std draft-ietf-6tisch-architecture-30*, pp. 1–60, Nov 2020.
- [9] M. R. Palattella, P. Thubert, X. Vilajosana, T. Watteyne, Q. Wang, and T. Engel, 6TiSCH Wireless Industrial Networks: Determinism Meets IPv6. Cham: Springer International Publishing, 2014, pp. 111–141.
- [10] X. Vilajosana, T. Watteyne, T. Chang, M. Vučinić, S. Duquennoy, and P. Thubert, "IETF 6TiSCH: A Tutorial," *IEEE Communications Surveys Tutorials*, vol. 22, no. 1, pp. 595–615, 2020.
- [11] S. Scanzio, G. Cena, A. Valenzano, and C. Zunino, "Energy Saving in TSCH Networks by Means of Proactive Reduction of Idle Listening," in Ad-Hoc, Mobile, and Wireless Networks, L. A. Grieco, G. Boggia, G. Piro, Y. Jararweh, and C. Campolo, Eds. Cham: Springer International Publishing, 2020, pp. 131–144.
- [12] G. Cena, S. Scanzio, L. Seno, A. Valenzano, and C. Zunino, "Energy-Efficient Link Capacity Overprovisioning In Time Slotted Channel Hopping Networks," in 16th IEEE International Conference on Factory Communication Systems (WFCS 2020), 2020, pp. 1–8.
- [13] I. Juc, O. Alphand, R. Guizzetti, M. Favre, and A. Duda, "Energy consumption and performance of IEEE 802.15.4e TSCH and DSME," in 2016 IEEE Wireless Communications and Networking Conference, 2016, pp. 1–7.
- [14] G. Cena, C. G. Demartini, M. Ghazi Vakili, S. Scanzio, A. Valenzano, and C. Zunino, "Evaluating and Modeling IEEE 802.15.4 TSCH Resilience against Wi-Fi Interference in New-Generation Highly-Dependable Wireless Sensor Networks," *Ad Hoc Networks*, vol. 106, p. 102199, Sep. 2020.