# Mitigating the Adverse Effects of Temperature on Low-Power Wireless Protocols

Carlo Alberto Boano and Kay Römer
Institute for Technical Informatics
Graz University of Technology, Austria
{cboano, roemer}@tugraz.at

Nicolas Tsiftes
SICS Swedish ICT
Stockholm, Sweden
nvt@sics.se

***Abstract*—Research and industrial installations have shown that the on-board temperature of wireless sensor nodes deployed outdoors can experience high fluctuations over time with a large variability across the network. These variations can have a strong impact on the efficiency of low-power radios and can significantly affect the operation of communication protocols, often compromising network connectivity. In this paper, we show the adverse effects of temperature on communication protocols and propose techniques to increase their resilience. First, we experimentally show that fluctuations of the on-board temperature of sensor nodes reduce the efficiency of data link layer protocols, leading to a substantial decrease in packet reception rate and to a considerable increase in energy consumption. Second, we investigate the reasons for such performance degradation, and show that high on-board temperatures reduce the effectiveness of clear channel assessment, compromising the ability of a node to avoid collisions and to successfully wake-up from low-power mode. After modelling the behaviour of radio transceivers as a function of temperature, we propose two mechanisms to dynamically adapt the clear channel assessment threshold to temperature changes, thus making data link layer protocols temperature-aware. An extensive experimental evaluation shows that our approaches considerably increase the performance of a network in the presence of temperature variations commonly found in real-world outdoor deployments, with up to 42% lower radio duty-cycle and 87% higher packet reception rate.***

***Keywords*—Clear Channel Assessment, CSMA Protocols, Outdoor Networks, Temperature Variations, Wireless Sensor Networks.**

## I. Introduction and Motivation

Temperature has a strong impact on the performance of wireless sensor networks. Real-world deployments have shown that the on-board temperature of wireless sensor nodes deployed outdoors can be significantly higher than air temperature [1]. Sensor nodes are indeed often exposed to direct sunlight and embedded into airtight packaging absorbing IR-radiation [2], causing the inner temperature in the casing to reach values as high as 70°C [3]. In a long-term outdoor deployment, Wennerström et al. [4] have indeed observed that the on-board temperature of a sensor node enclosed into an airtight packaging can experience variations up to 83°C across different seasons, and 56°C within 24-hours [5], with large heterogeneity across the network [6].

These temperature fluctuations can have a strong impact on clock drift, slowing down processor operations [6] and affecting time synchronization between nodes [7]; as well as on the lifetime of sensor nodes, influencing the capacity and discharge curve of batteries [8], [9] and altering the current consumption of electronic components [10], [11].
Furthermore, temperature can also drastically affect the efficiency of low-power wireless transceivers and reduce the quality of wireless links. The performance of low-power radios employed in off-the-shelf wireless sensor nodes is indeed temperature-dependent [12], with a reduction in the strength of the transmitted and received signal at high temperatures. For example, a temperature variation of 40°C can decrease the strength of the received signal by up to 6 dB, with a negative effect on the correct reception of packets [5].

To better study the impact of temperature variations on low-power wireless links and communications protocols, we have designed TempLab, a testbed infrastructure with the ability of varying the on-board temperature of sensor nodes and reproducing the temperature fluctuations found in outdoor deployments [6]. We have shown how this temperature-controlled testbed can be used to systematically analyse the performance of communication protocols, and highlighted that the latter exhibit a substantially lower efficiency at high temperatures.

In this paper, we exploit this temperature-controlled testbed to analyse in detail the performance of state-of-the-art communication protocols and to understand (i) *why* their performance decreases in the presence of temperature variations, and (ii) *how* we can mitigate the problem and improve their resilience towards temperature fluctuations. We first show experimentally that fluctuations of the on-board temperature of sensor nodes reduce the efficiency of carrier sense multiple access data link layer protocols, leading to a substantial decrease in the packet reception rate and to an increase of the energy consumption. We identify reduced effectiveness of clear channel assessment as the reason for such performance degradation, and show that this reduced effectiveness compromises the ability of a node to avoid collisions and to successfully wake-up from low-power mode. Based on these insights, we propose two mechanisms to mitigate the problem by dynamically adapting the clear channel assessment threshold to temperature changes: one based on the temperature measured locally, and one on the highest temperature measured across all neighbouring nodes. We finally show through an extensive experimental evaluation that the proposed approaches increase the robustness of existing protocols to temperature variations and significantly improve the performance also on a network level.

The contributions of this paper are hence three-fold:

- **Inefficiency of clear channel assessment.** We describe how temperature variations affect the efficiency of clear channel assessment, and show experimentally that this inefficiency compromises the operations of data link layer protocols based on carrier sense.
- **Adaptive data link layer protocols.** After modelling the behaviour of radio transceivers as a function of temperature, we implement two strategies that increase

the efficiency of clear channel assessment by making data link layer protocols temperature-aware.

- **Extensive experimental evaluation.** We show that our improved protocols sustain a significantly higher performance than existing protocols, with up to 71% lower energy consumption and 194% higher packet reception rate in the presence of temperature variations commonly found in real-world outdoor deployments.

The next section describes the impact of temperature on low-power radios, and models the attenuation of signal strength on the platform used in our experiments. Sect. III analyses the impact of temperature on data link layer protocols, and highlights the inefficiency of clear channel assessment at high temperatures. In Sect. IV we describe two mechanisms to correct this inefficiency and to make data link layer protocols temperature-aware. We evaluate the performance of our approaches in Sect. V, showing large performance improvements on a link basis and on a network level. After describing related work in Sect. VI, we conclude our paper in Sect. VII.

## II. Impact of Temperature on Low-Power Radios

Experiences and reports from long-term outdoor deployments have highlighted that temperature has a strong impact on the performance of low-power radio transceivers.

*Impact of temperature on link quality.* Results by Bannister et al. [12] from an outdoor deployment in the Sonoran desert have revealed that an increase in temperature causes a reduction of the wireless link quality. These results were later confirmed by indoor and outdoor experiments [2], [3], and by a long-term outdoor deployment by Wennerström et al. [4] in Uppsala, Sweden. In the latter, 16 TelosB nodes equipped with the CC2420 radio were placed within each other's transmission range, and exchanged packets and recorded statistics for several months. Fig. 1 shows the data collected by two nodes in this deployment: the top figure shows the temperatures measured on-board and the air temperature recorded by a nearby weather station; the other figures show the evolution of a number of link quality metrics over time. Firstly, we can observe that the on-board temperature of the sensor nodes is significantly higher than air temperature: this is very common in outdoor deployments when nodes are enclosed into airtight packaging absorbing IR-radiation. Secondly, we can observe a clear correlation between the on-board temperature of the two nodes and the quality of their link: the higher the temperature, the lower the received signal strength indicator (RSSI) and the link quality indicator representing the chip error rate (LQI).

*Dependency between temperature and signal strength.* Bannister et al. [12] have shown that the attenuation in received signal strength on the CC2420 radio chip is the result of the decreased efficiency of the transmitter's power amplifier and the receiver's low-noise amplifier at high temperatures. In their experiments in a climate chamber, the authors observed a decrease of 4-5 dB in the output power of the transmitter and a drop of 3-4 dB in the received power over the temperature range 25-65 °C, for a combined effect on received signal strength of 8 dB when both transmitter and receiver are heated. We have confirmed in later experiments over a larger temperature range [5] that the relationship between temperature and signal strength attenuation is approximately linear, and that this also applies to other radio chips employed in off-the-shelf sensornet platforms. Fig. 2 shows the strength of the received signal at different temperatures between two Maxfor MTM-CM5000MSP sensor nodes (replica of TelosB motes) while the transmitter, receiver, or both transmitter and receiver nodes are heated using TempLab [6]. We can notice that the received signal strength attenuation is similar when the two nodes are heated individually (a loss of 0.08 dB/°C[1]), and about twice as high when both nodes are heated at the same time (a loss of 0.17 dB/°C). Instead, the noise floor, i.e., the received signal strength measured in absence of radio activity, exhibits a lower variability in the presence of temperature variations.

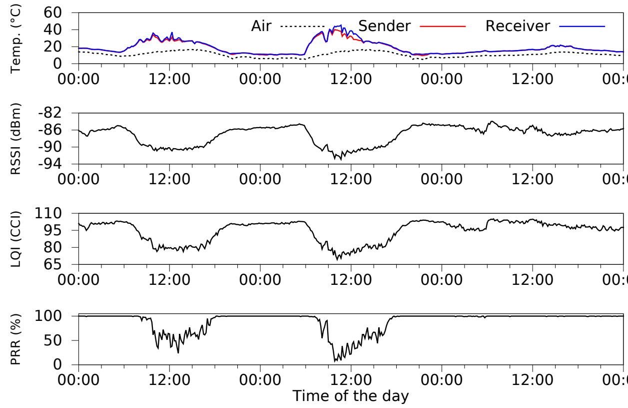


Fig. 1. High temperatures decrease the performance of low-power radios. In traces from Wennerström et al.'s outdoor deployment [4], we can observe that during daytime (when temperature is high), the received signal strength indicator (RSSI) and link quality indicator (LQI) are lower than during the night. During daytime, also the packet reception rate (PRR) is reduced.

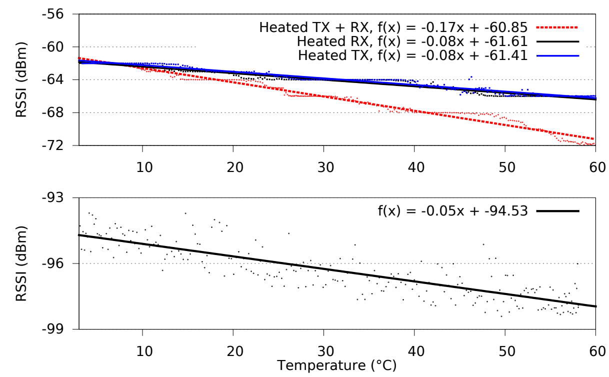


Fig. 2. Signal strength attenuation as a function of temperature. The top plot shows the received signal strength of packets while transmitter (blue), receiver (black), or both transmitter and receiver (red) are heated: the attenuation is highest when both nodes are heated at the same time. The bottom plot shows the received signal strength attenuation in absence of packet transmissions.

*Impact on packet reception.* The attenuation of the signal strength at high temperatures can affect the reception of packets in two ways. First, a weaker signal is more susceptible to bursts of external interference, and the probability that devices operating at higher powers (e.g., Wi-Fi access points and microwave ovens) corrupt or destroy a packet increases at high temperatures. Second, if temperature increases and the signal strength weakens to values close to the ambient RF noise (often called noise floor), the radio's ability to successfully demodulate a packet significantly decreases. When

[1] We estimate the attenuation by computing the slopes of the RSSI curve. Please note that an *exact* comparison between two curves is not possible, as RSSI readings are integer values that depend on the operation of the automatic gain controller and on the hysteresis between different gain modes [5].

this happens, a physical limit is reached: the radio cannot correctly receive (most of) the packets that were transmitted, and the connectivity of the link is irreparably compromised. This situation is captured in Fig. 1 (bottom). In Wennerström et al.'s deployment, the nodes communicate using Contiki's nullMAC, a data link layer protocol in which the radio remains active all the time and packets are transmitted without first verifying the absence of other traffic. As soon as the received signal strength weakens to values close to the noise floor in the deployment environment ($\approx$ -94 dBm), the packet reception rate (PRR) between the two nodes drops significantly, and the link becomes almost useless during daytime.

In the next section, we focus on carrier sense multiple access data link layer protocols and show that their performance decreases significantly at high temperatures, *but not as a result of the above observations*. The vast majority of duty-cycled MAC protocols *do not* actually reach the physical limit of the radio at high temperatures, and the lower reception rates are caused by design choices that neglect the inefficiency of clear channel assessment in the presence of temperature fluctuations.

## III. Impact of Temperature on CSMA Protocols

The attenuation of received signal strength at high temperatures described in Sect. II can affect two key functionalities of carrier sense multiple access (CSMA) protocols.

1) *Collision avoidance.* CSMA protocols rely on clear channel assessment (CCA) to determine whether another device is already transmitting on the same frequency channel, and defer transmissions that may otherwise collide with ongoing communications.
2) *Wake-up of nodes.* Duty-cycled protocols typically employ CCA to trigger wake-ups, i.e., to determine if a node should stay awake to receive a packet or whether it should remain in low-power mode.

CCA implementations are typically based on energy detection, i.e., on the measurement of the received signal strength and on its comparison with a given threshold. When performing energy detection using a fixed CCA threshold, it is neglected that received signal strength readings are affected by temperature, and this leads to a number of problems. First, the transmitter can erroneously measure a weaker noise in the environment as a result of the increased temperature, and generate wasteful transmissions (see Sect. III-B). Second, a receiver node may not receive a signal sufficiently strong to cause a wake-up of the radio, and constantly remain in low-power mode at high temperatures, causing the disruption of the link (see Sect. III-C). We analyse these issues in the remainder of this section, after describing how CCA is typically implemented in sensornet MAC protocols.

### A. Clear Channel Assessment in Sensornet MAC Protocols

In CSMA protocols, the correct operation of clear channel assessment is fundamental to reduce the number of wasteful transmissions and to preserve the limited energy budget of the nodes in the network. The typical task of CCA is to avoid collisions, i.e., to determine whether another device is already transmitting on the same frequency channel. If there are ongoing transmissions, CSMA protocols defer transmissions using different back-off strategies [13]; otherwise the packet(s) are immediately sent. CCA is also used in low-power duty-cycled MAC protocols to trigger wake-ups, i.e., to determine if a node should remain awake to receive a packet or whether it should return in sleep mode [14]. Towards this goal, low-power MAC protocols typically perform an inexpensive CCA check and keep the transceiver on if some ongoing activity is detected on the channel [14], [15], [16].

The CCA check can be carried out using energy detection or carrier sense, as described in the IEEE 802.15.4 standard. Energy detection consists in sampling the energy level in the wireless channel and determining whether another device is already transmitting by comparing the measured signal strength with a given CCA threshold $T_{CCA}$. Carrier sense consists in detecting the presence of a modulated signal, irrespective of its strength. Both options can also be used at the same time: in the CC2420 transceiver, this is the default configuration.

*Most protocols employ fixed CCA thresholds.* When using energy detection, a critical design choice is the selection of $T_{CCA}$. Whilst sender-initiated, duty-cycling MAC protocols such as B-MAC [14], BoX-MACs [17], and ContikiMAC [16] include energy detection as an important feature to reduce idle listening, there is not yet a widespread practice of tuning the CCA threshold at run-time in relation to the noise floor of each network deployment. Rather, the current practice is to rely on the default system settings, i.e., on a *fixed* CCA threshold, which is either set at compile-time, or left untouched so that the default value of the radio device is used instead. The IEEE 802.15.4 standard suggests to use a $T_{CCA}$ that is at most 10 dB greater than the radio's specified receiver sensitivity. Contiki uses the default value for most hardware platforms (the CC2420's default threshold is -77 dBm), but did recently set $T_{CCA}$ for TelosB-based platforms to -90 dBm.

### B. Inefficient Collision Avoidance

When a protocol employs a fixed CCA threshold to determine whether another device is already transmitting, it essentially neglects that the received signal strength depends on the temperature. We now show experimentally that this can lead to an increase in false negatives when a transmitter is assessing the presence of a busy medium.

Fig. 3(a) shows an overview of our testbed, equipped with eighteen Maxfor MTM-CM5000MSP nodes. We use TempLab [6] to vary the on-board temperature of the nodes between 25 and 75°C using IR heating lamps (Fig. 3(b)). We carry out experiments consisting of several transmitter-receiver pairs running a basic Contiki application, in which the transmitter node periodically sends packets to its intended receiver and collects statistics such as the energy expenditure at the link-layer and the RF ambient noise in the radio channel. The latter is computed as the maximum of 20 consecutive RSSI readings after a packet transmission. In a first experiment in an environment rich of Wi-Fi interference, we use Contiki's nullMAC and nullRDC to avoid protocol-specific implementations and employ the CC2420's default CCA threshold (-77 dBm). Except from temperature, there is no significant change in the environment surrounding the nodes.

Fig. 3(c) shows the ambient noise captured using RSSI readings by a node in our testbed. The noise has a visible correlation with the on-board temperature of the node, and follows the attenuation described in Sect. II. We can observe

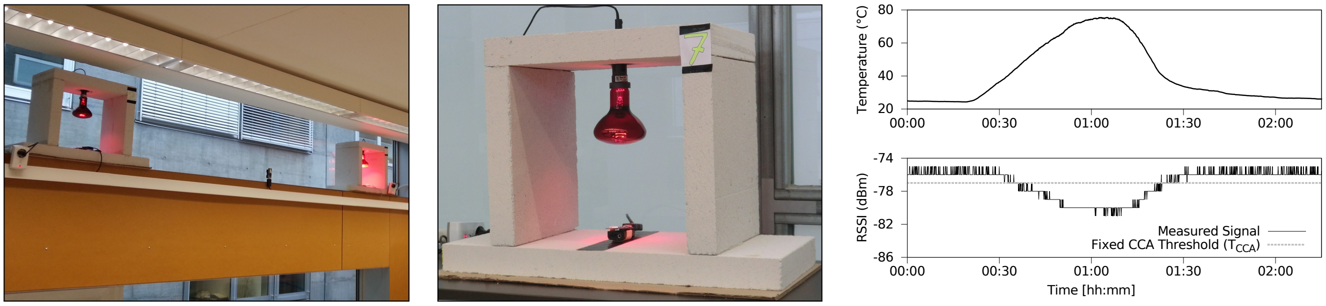


(a) Overview of our testbed infrastructure (b) IR heating lamp on top of a sensor node (c) Signal strength measured at high temperatures

Fig. 3. Overview of the testbed infrastructure used in our experiments (a) with infra-red heating lamps on top of each sensor node to control their on-board temperature (b). The received signal strength weakens at high temperatures and can cause an intersection with $T_{CCA}$, causing several issues (c).

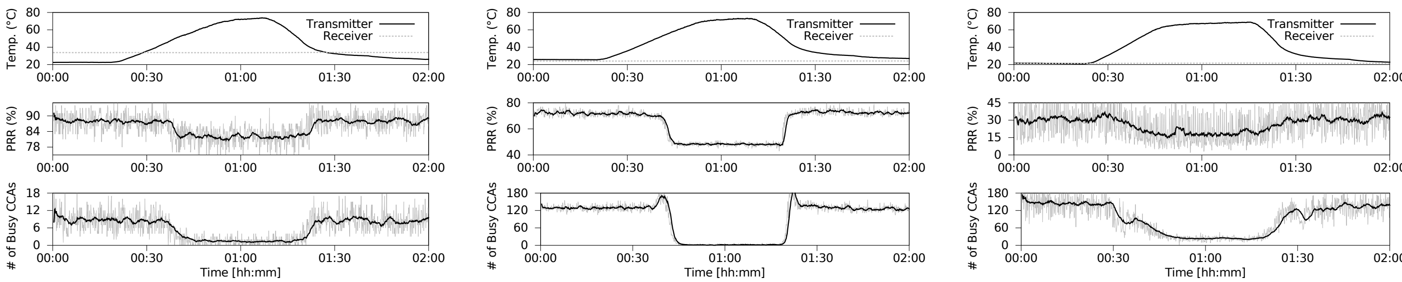


(a) JamLab's emulated Wi-Fi video streaming (b) JamLab's emulated microwave oven (c) File transfer between two Wi-Fi devices

Fig. 4. Temperature affects the efficiency of collision avoidance in CSMA protocols. Our experiments in different interference scenarios show that when the received signal strength weakens to values below $T_{CCA}$ at high temperatures, the PRR decreases, as well as the number of CCAs identifying a busy channel.

that at around 40°C, there is an intersection between the measured signal strength and the selected $T_{CCA}$. For temperatures lower than 40°C the measured RSSI is above $T_{CCA}$ (and hence transmissions would be deferred); for temperatures higher than 40°C, instead, the RSSI is below $T_{CCA}$ (and packets would be immediately sent). In other words, the MAC protocol erroneously deduces from RSSI readings obtained above 40°C that the channel is free from harmful interference. In reality, the interference in the environment is not weakened by temperature (the RSSI attenuation is only an artefact of the radio), and can still destroy transmitted packets. These erroneous clear channel assessments at high temperature may hence lead to an increase in the number of wasteful transmissions destroyed or corrupted by surrounding interference.

Fig. 4 shows the impact of erroneous clear channel assessments in the presence of different interference patterns. We use JamLab [18] to produce repeatable interference in our testbed on different channels. We emulate on one channel the interference caused by a computer streaming videos from a Wi-Fi access point, and on another channel the one caused by an active microwave oven. We also let a computer transfer large files from a nearby Wi-Fi access point using a channel that is not affected by JamLab. We then analyse how this affects the PRR on the transmitter-receiver pairs in our testbed that experienced an intersection between measured noise and $T_{CCA}$ at different temperatures as in Fig. 3(c). We can notice that in all scenarios the PRR decreases as soon as the on-board temperature of sensor nodes increases. In the presence of Wi-Fi video streaming, the PRR of the link decreases from 88 to 81% (Fig. 4(a)), whereas in the presence of an active microwave oven the PRR decreases from 70 to 45% (Fig. 4(b)). Similarly, also the PRR in the presence of a Wi-Fi file transfer decreases from 30 to 18% at high temperatures (Fig. 4(c)). We can also notice that the decrease in PRR is correlated with a decrease in the number of clear channel assessments identifying a busy channel, i.e., with a decrease in the number of clear channel assessments that do not identify potential collisions at high temperatures. These results prove our hypothesis, and show that the intersection between the RSSI curve and the CCA threshold shown in Fig. 3(c) results in erroneous clear channel assessments leading to a decreased PRR at high temperatures.

### C. *Unsuccessful Wake-Up of Nodes*

State-of-the-art MAC protocols often duty cycle the radio to reduce energy consumption, and employ clear channel assessment to wake-up the transceiver from sleep mode. Typically, a periodic CCA check is performed: if the channel is busy, the transceiver is kept on in order to receive the incoming packet, otherwise the radio returns to sleep mode.

High temperatures can affect the correctness of this mechanism. Imagine a sender $A$ and a receiver $B$ exchanging packets using a duty-cycled MAC protocol in which $A$ sends short strobes before the actual packet (or repeatedly sends the same packet). If $B$ receives the strobes from node $A$ with a signal strength that is higher than $T_{CCA}$, it keeps its radio on and receives the payload message from $A$. If temperature increases, the received signal strength at node $B$ may intersect $T_{CCA}$ as shown in Fig. 3(c). When this happens, the transmissions from $A$ are received with a signal strength lower than $T_{CCA}$, and $B$ does not wake up to receive $A$'s packets anymore, essentially disrupting the link. In the case shown in Fig. 3(c), the link would be disrupted for temperatures higher than 40°C, because node $B$ would not wake up when the strength of the received signal from $A$ decreases below $T_{CCA}$.

Please note that the probability that the received signal strength intersects $T_{CCA}$ as a result of an increase in temperature can be quite high. Temperature variations can cause a

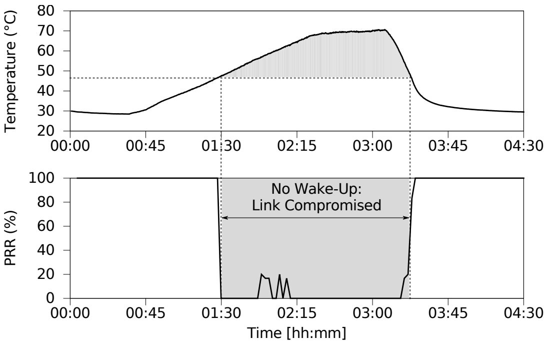


Fig. 5. Temperature can affect the wake-up mechanism in duty-cycled MAC protocols. When the strength of the received signal from a transmitter weakens at high temperatures and intersect the CCA threshold as shown in Fig. 3(c), the receiver does not wake up anymore, disrupting the link's connectivity.

signal attenuation by up to 10 dB (as shown in Fig. 2), which implies that all links in a network with an RSSI between $T_{CCA}$ and $10+T_{CCA}$ are prone to this problem. For example, when using the CC2420 radio ($T_{CCA}$=$-77$ dBm) and transmitting at 0 dBm, the majority of nodes with a distance between 5 and 25 meters would form a link with an RSSI falling in this range [19].

We now show experimental evidence of this problem. We let several transmitter-receiver pairs of nodes communicate using ContikiMAC, Contiki's default MAC protocol in which nodes sleep most of the time and periodically wake up to check for radio activity. In ContikiMAC, the transmitter sends repeatedly the same packet until a link layer acknowledgement (ACK) is received, whereas the receiver keeps its radio on as soon as a packet transmission is detected by means of a single CCA check [16]. Packets are exchanged every 20 seconds, and ACKs are sent using CC2420's hardware support. As in the previous experiment, we use TempLab to warm-up and cool-down the on-board temperature of the nodes, emulating the daily fluctuations that can be found in real-world deployments.

Fig. 5 shows an example of link disruption caused by a receiver not waking up at high temperatures. We can notice that what was a perfect link until approximately 47°C, suddenly does not receive any packet at higher temperatures. Only once temperature decreases below 47°C, the link is restored and the node correctly receives the packets sent from the transmitter. This behaviour can significantly harm network performance, as links may disappear during the hottest times of the day, leading to high latencies, drastic topology changes, or in case no alternative paths for communication can be found, to a complete disconnection of some nodes from the network.

A receiver node could in principle detect a packet transmission using carrier sense, i.e., by identifying a valid sequence of bits without comparing if the received energy is above a given threshold. However, in off-the-shelf radio transceivers such as the CC2420, a valid sequence can be identified only prior detection and validation of the start frame delimiter. Therefore, carrier sense is ineffective when used in duty-cycled systems that periodically wake up and perform a single CCA check (in a non-duty-cycled protocol such as Contiki's nullRDC or nullMAC, instead, carrier sense would work well, as the radio remains always active). Indeed, despite the CC2420 radio uses by default a combination of carrier sense and energy detection, ContikiMAC experiences a complete loss at high temperatures that is dependent on the chosen energy detection threshold.

It is also important to highlight that selecting by default a low CCA threshold is not optimal: the lower $T_{CCA}$ the higher the number of activities in the channel (radio interference, communications from surrounding nodes) that will trigger a wake-up and, consequently, a higher energy consumption. Indeed, selecting $T_{CCA}$ close to the noise floor in a noisy environment, would essentially cause the radio to be almost constantly active, with a highly suboptimal energy expenditure.

## IV. DESIGNING TEMPERATURE-AWARE MAC PROTOCOLS

Whenever a link delivers poor performance, it is typically the network layer's task to maintain connectivity and look for alternative routes that can sustain a high delivery rate. Using link quality estimation, the network layer can indeed filter out lossy links and pick a better topology, i.e., select a network configuration that avoids links that are asymmetric or that have a signal that is too weak to communicate reliably, as well as links that are negatively affected by temperature variations. The network layer, however, can only be effective if the network is sufficiently dense to offer a high link redundancy: very often there are no available neighbours forming a link that can offer a better performance, especially in sparse networks. In such cases, the network layer is obliged to make use of lossy links, and cannot mitigate the impact of temperature variations on the lower layers of the protocol stack.

To mitigate the inefficiency of CSMA protocols at high temperatures shown in Sect. III, we hence need to tackle the problem directly at the MAC layer. A link can indeed still offer good performance if the CCA threshold is dynamically adapted to the on-board temperature variations of the nodes. In this section, we propose two alternatives to achieve this goal.

### A. Predicting the Attenuation of Signal Strength

In order to dynamically adapt $T_{CCA}$ to temperature variations, we first need to model the relationship between signal strength attenuation and temperature. In Sect. II we have shown that the latter is approximately linear, and that there are two components that need to be considered: the attenuation on the receiver side, and the one on the transmitter side.
Imagine a sender $A$ and a receiver $B$ exchanging packets. If the on-board temperature of $B$ varies by $\Delta T_B$ degrees w.r.t. to an initial temperature $\tau$, the signal will suffer an attenuation on the receiver side by $R = \beta \Delta T_B$, with $\Delta T_B$ being the difference between $B$'s current temperature $T_{now}$ and $\tau$. Similarly, if the on-board temperature of $A$ varies, its signals will be transmitted with an attenuation on the transmitter side of $T = \alpha \Delta T_A$, and $B$ will receive a signal that is $T$ dBm weaker. In case the temperatures of both $A$ and $B$ vary, the overall attenuation of the received signal strength on $B$ is given by $R + T$. Please notice that if temperature has decreased, $\Delta T = (T_{now} - \tau)$ is negative, and $R$ and $T$ are not an attenuation, but instead a strengthening of the signal.

$\alpha$ and $\beta$ are specific to the employed radio and differ only in a negligible way among different instances of the same chip. Hence, they can be characterized following the same approach shown in Sect. II: using a pair of nodes that can be heated individually, we compute the variation of signal strength on a large temperature range and derive the slope of the RSSI curves of transmitter and receiver for a given platform [5].

In the case of the Maxfor nodes employed in our experiments we derive from Fig. 2 $\alpha$ = $\beta$ = -0.08 dB/°C. We further model the attenuation of the noise floor as $N = \gamma\Delta_T$ (which is typically smaller than $R$ and $T$) and derive $\gamma$ = -0.05 dB/°C.

### B. Adapting the CCA Threshold at Runtime

Exploiting the above model, we can now adapt the CCA threshold at runtime. Each node needs to compute if temperature varied significantly enough to cause an attenuation of the signal strength w.r.t. an initial threshold $T'_{CCA}$.
As we mentioned in Sect. III, the default CCA threshold is typically fixed. However, as nodes are typically uncalibrated and have radio irregularities, a good practice would be to select $T'_{CCA} = n'_f + K$, with $n'_f$ being the noise floor of the node, and $K$ a constant defined at compile time. If this is the case, $T'_{CCA}$ and $n'_f$ are computed during the start-up phase while the node experiences an on-board temperature $\tau$. If $T'_{CCA}$ is fixed, we assume $\tau = 25$°C. Please note that high values of $K$ reduce the number of activities in the channel that can trigger a wake-up of a node (minimizing energy consumption), but also reduce the number of links in the network (fewer neighbours can wake-up a node with a signal strength higher than $T'_{CCA}$).
Whenever temperature varies significantly, we compute the updated threshold as $T_{CCA} = T'_{CCA} + T + R$, with $T$ and $R$ being computed using the difference between the current temperature and $\tau$. We apply to the computation of $T_{CCA}$ a lower bound $n_f + C$ (with $n_f = n'_f + N$) that avoids the selection of CCA thresholds that are too close to the noise floor (this would cause the radio to continuously wake-up).

*Obtaining up-to-date temperature measurements.* All that is needed to adapt the threshold is hence an up-to-date information about the current on-board temperature of the nodes and the initial temperature $\tau$ stored in a 2-byte variable. Almost every off-the-shelf sensornet platform comes with an embedded temperature sensor. TelosB-based platforms carry the SHT11, a digital temperature and humidity sensor. Other platforms do not have a dedicated sensor, but several microcontrollers such as the MSP430 offer the possibility to obtain a rough estimate of the on-board temperature from a built-in temperature sensor using a specific input of the analog-to-digital converter. By periodically sampling the on-board temperature, a node can hence compare its current temperature with $\tau$ and compute $\Delta_T$. It is important to stress that the temperature sensor should be physically on the board, to get an estimate as close as possible to the temperature of the radio chip: external sensors measuring air temperature outside the packaging may not give a sufficiently accurate estimation.

*Deriving the on-board temperature of the transmitter.* By knowing its current on-board temperature, a node can immediately derive $N$ and $R$. If a node would adapt its CCA threshold based on this information (i.e., using $T = 0$), the inefficient collision avoidance problem at high temperatures would be solved, as well as the wake-up problem in case the temperature of the transmitter node does not vary significantly. If this node, however, receives packets sent from a node experiencing temperature fluctuations, it would need to know the temperature of the transmitter to derive $T$ and completely mitigate the unsuccessful wake-up problem. This is a non-trivial problem, as a receiver does not necessarily know the identity of the sender by the point in time in which it performs a CCA, and as it may actually be recipient of packets sent by different nodes. Assuming that transmitter and receiver experience the same temperature variations may lead to inaccurate results: real-world deployments have shown that there can be high gradients (more than 30°C) even across spatially close nodes [6], [7] because of cloud obstructions or shade from trees or buildings in the surroundings. Similarly, setting a fixed worst-case temperature at compile-time would lead to suboptimal performance, as $T_{CCA}$ would remain unnecessarily low most of the time.
The information about the transmitter's temperature can actually be conveyed by the network layer, which stores a table of neighbour addresses and attributes, and can be augmented with an attribute for the latest on-board temperature of each neighbour. Modifying the network layer in this manner may not be suitable in all systems, however. Hence, we propose two different adaptation mechanisms: one that adapts $T_{CCA}$ based only on local temperature measurements, and one that exploits a cross-layer approach to derive $T$.

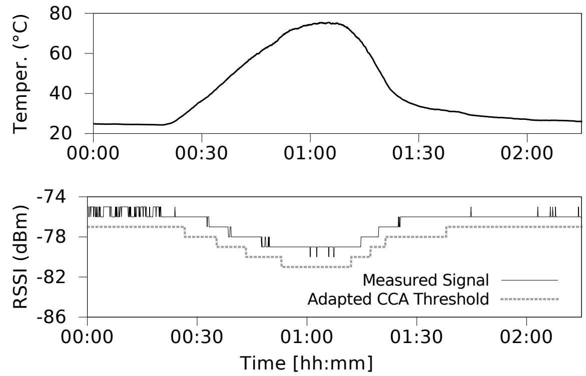


Fig. 6. Dynamic adaptation of the CCA threshold based on the temperature measured locally on the node: $T_{CCA}$ follows the attenuation of the signal, avoiding an intersection with the RSSI curve (in contrast with Fig. 3(c)).

*Local adaptation.* A first approach adapts $T_{CCA}$ based on local temperature measurements only (i.e., it fixes $T$=0). In this case, $T_{CCA} = T'_{CCA} + R$, with a lower bound $n_f + C$. We found in our experiments that values of $C$ below 2 dBm trigger an almost continuous wake-up of the radio, and we therefore use $C$=2 dBm. Fig. 6 shows the adaptation of the CCA threshold based on the algorithm detailed previously. We replicate the setup of Sect. III-B and heat a receiver node while measuring the strength of the signal in an environment rich of Wi-Fi interference. If we compare the results with the ones shown in Fig. 3(c), we can notice that the CCA threshold follows the same attenuation as the received signal, avoiding an intersection between the RSSI curve and $T_{CCA}$. This shows that the proposed model is sufficiently accurate to dynamically adapt $T_{CCA}$ to local temperature changes. However, if the on-board temperature of the transmitter significantly varies, a receiver node may still experience unsuccessful wake-ups.

*Cross-layer adaptation.* To prevent this, we propose an approach that allows the CCA adaptation mechanism to make more informed decisions by using temperature information from the neighbours. Our cross-layer adaptation uses existing routing beacons to piggyback temperature information efficiently, and computes the maximum temperature change across all neighbours. We implement this by using RPL, the standard IPv6 routing protocol for low-power and lossy networks [20].

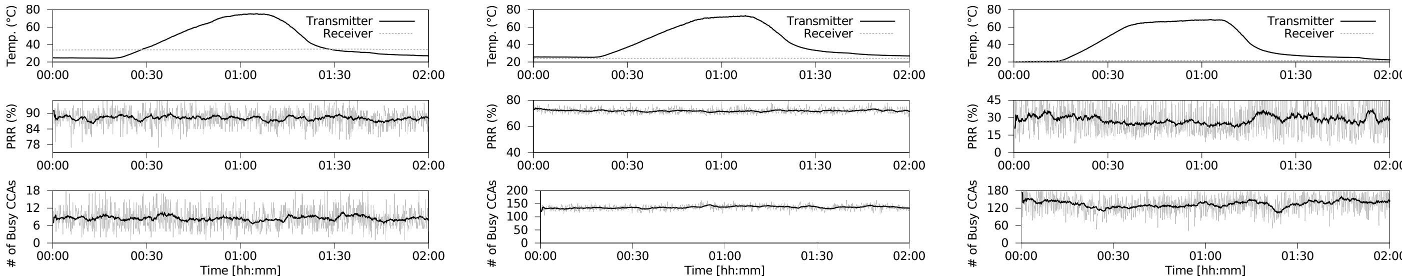


(a) JamLab's Wi-Fi Video Streaming (b) JamLab's Wi-Fi Active Microwave Oven (c) Real Wi-Fi File Transfer

Fig. 7. When adapting the CCA threshold based on local temperature measurements, temperature does not affect the efficiency of collision avoidance in CSMA protocols. In contrast with the results shown in Fig. 4, the PRR remains fairly constant for all interference scenarios despite temperature variations.

Whilst we have chosen RPL because it is a standard protocol and several open-source implementations exist, we also note that it would be simple to disseminate the information at the application layer, albeit with a slightly higher energy cost. We disseminate the temperature information by piggybacking it on RPL's routing beacons. RPL sends these beacons to the neighbour nodes with quickly increasing time intervals, as regulated by the Trickle algorithm [21]. Within the DODAG Information Object (DIO), there is room to embed a routing metric container object, which holds different parameters and constraints that are used to take routing decisions. Beside the metric container specified in the standard, it is possible to use implementation-defined metric containers. Hence, we make each node report its current and maximum temperature through such a metric container. Once a node receives this information in an incoming routing beacon, it stores it as an attribute in Contiki's neighbour table, from whence it can be retrieved by the CCA adaptation module to calculate the maximum temperature change in the neighbourhood.

## V. Evaluation

We now evaluate the performance of our approaches experimentally. We first show that they alleviate the collision avoidance and wake-up problem in CSMA protocols. We then run a network of nodes, and show that when employing a MAC protocol with an adaptive threshold, the performance of the network significantly increases, with up to 42% lower radio duty cycle and 87% higher PRR in the presence of temperature variations commonly found in outdoor deployments.

### A. Improved Collision Avoidance

In Sect. III-B we have shown that with varying on-board temperatures, a transmitter can erroneously measure a weaker noise and generate wasteful transmissions. Using the same experimental setup, we now analyse the performance of the transmitter-receiver pairs in our testbed when dynamically adapting $T_{CCA}$ using local temperature information. We use the CC2420's default CCA threshold, i.e., $T^{'}_{CCA} = -77$ dBm and use Contiki's nullMAC and nullRDC. Fig. 7 shows the PRR experienced by the links in the same interference scenarios described in Sect. III-B (the experiments were executed back-to-back). If we compare the results with Fig. 4, we can notice that the PRR does not depend on the on-board temperature of the nodes, but remains instead fairly constant throughout the experiment. This hints that the adapted protocol is able to avoid the intersection between the RSSI curve and $T_{CCA}$, mitigating the collision avoidance problem.

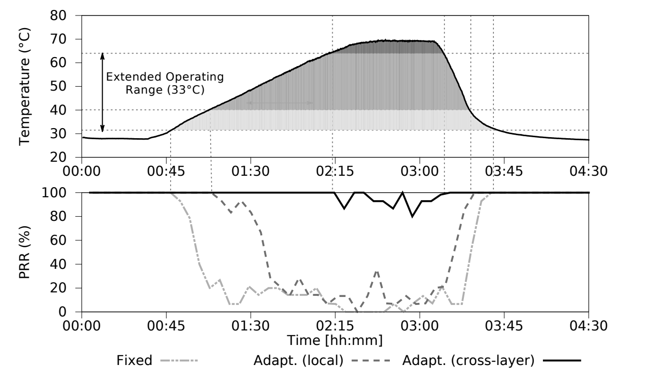


Fig. 8. Adaptive CCA thresholds alleviate significantly the wake-up problem at high temperatures. By adapting $T_{CCA}$, we can extend the usability of a link at much higher temperatures.

### B. Improved Wake-Up Efficiency

In Sect. III-C we have shown that a receiver node exposed to temperature variations may not receive a signal sufficiently strong to cause a wake-up of the radio, and constantly remains in low-power mode, causing the disruption of the link. We employ ContikiMAC with a $T^{'}_{CCA} = n^{'}_{f} + K$ with $K = 6$ dBm and use TempLab to warm-up and cool-down the on-board temperature of both transmitter and receiver, emulating the daily fluctuations that can be found in real-world deployments. We repeat the experiments several times and run (i) an unmodified ContikiMAC using a fixed CCA threshold, (ii) an adaptive threshold based on local temperature information, and (iii) an adaptive threshold based on the information inferred from the routing layer. Fig. 8 shows the PRR on a representative link in our testbed (a similar trend was observed across all links): the adaptation of the CCA threshold can significantly alleviate the wake-up problem. When using a fixed threshold, the link starts to experience packet loss at 31°C. Instead, the link sustains 100% delivery rate up to 40°C when using local temperature information and up to 64°C when using the information inferred from the routing layer. This essentially implies that the use of a dynamic $T_{CCA}$ extends the usability of a link to a higher temperature. It is important to highlight that the adaptation of $T_{CCA}$ does not mitigate completely the impact of temperature. The reason lies in the selection of $T^{'}_{CCA}$: by selecting $K = 6$ dBm, the high temperature variation attenuates the signal strength by several dB, reaching the physical limit of the radio (i.e., at temperatures higher than 64°C we receive a signal strength that is too weak to be successfully demodulated). Hence, the higher is $K$, the higher can be the performance gain compared to a protocol using a fixed CCA threshold.

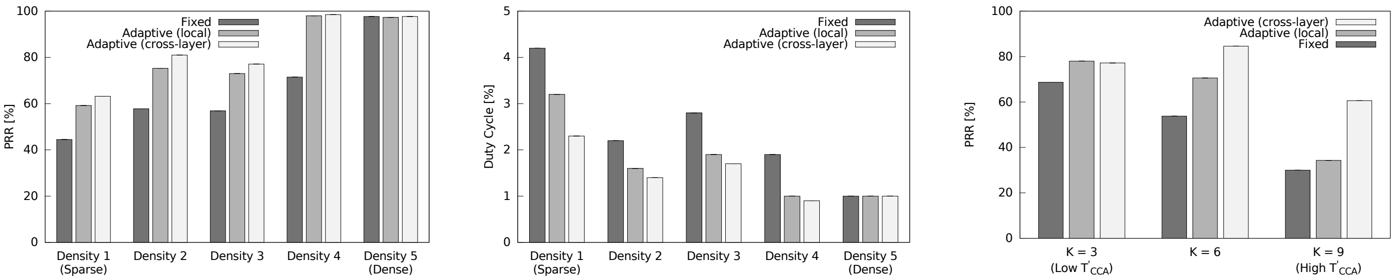


(a) PRR as a function of network density (b) Duty cycle as a function of network density (c) PRR as a function of $T^{'}_{CCA}$

Fig. 9. Network performance when using fixed and adaptive CCA thresholds as a function of network density (a, b), and $T^{'}_{CCA}$ (c).

## C. Performance on a Network Level

We now present results obtained running a data-collection protocol on several networks, and show the benefits of using dynamically adapted CCA thresholds in the presence of temperature variations. We use RPL in our testbed deployed in a 55 $m^2$ room: we select one node as a sink, and we create five different network densities by using only a portion of the nodes: 5, 7, 9, 11, and 13 nodes, respectively. By varying the density from roughly one node every 11 $m^2$ to a node every 4 $m^2$, we can see largely different impacts on a network level, as the ability of the network layer to select alternative links is constrained. Using the same temperature profiles and setup as in the previous example (all nodes use a transmission power of -25 dBm), we carry out experiments with ContikiMAC using: (i) a fixed CCA threshold, (ii) an adaptive threshold based on local temperature information, and (iii) an adaptive threshold using the information inferred from the network layer.

Our results indicate that temperature strongly affects network performance, especially in sparse networks. Fig. 9(a) shows that if the network is dense, the routing layer can mitigate the impact of temperature and sustain a high PRR even with a MAC protocol employing a fixed CCA threshold. The less dense the network is, the higher becomes the impact of temperature on a protocol using a fixed threshold, with the average PRR in the network dropping below 50%. Instead, when using adaptive thresholds, the network sustains higher reception rates in sparse networks (from 44 to 63%, and from 57 to 81% in the two sparsest configurations), with the highest PRR recorded when using the information inferred from the routing layer in line with the experiments in Sect. V-B.

We further analyse the energy-efficiency of the different approaches by comparing the average radio duty cycle in the network. Fig. 9(b) shows that adaptive CCA thresholds sustain significantly lower duty cycles, as a result of a reduced number of retransmission attempts and wasteful transmissions. In the sparsest network configuration, the duty cycle drops from 4.2% to 3.2% in the case of local temperature information and to 2.3% when using the temperature inferred from the routing layer. The latter corresponds to a 55% higher energy-efficiency than when using a fixed threshold. With denser networks the duty cycle decreases, as the network layer can select alternative links and seamlessly mitigate the impact of temperature.

Fig. 9(c) shows the role of the initial CCA threshold $T^{'}_{CCA}$ in a network with a density of one node every 8 $m^2$. We set $T^{'}_{CCA} = n^{'}_{f} + K$ using different $K$ values, and show that the higher $K$ is, the higher are the performance improvements introduced by the adaptive approaches. This is the result of the observation made in Sect. V-B: the higher $K$ is, the more the usability of a link can be extended at high temperatures.

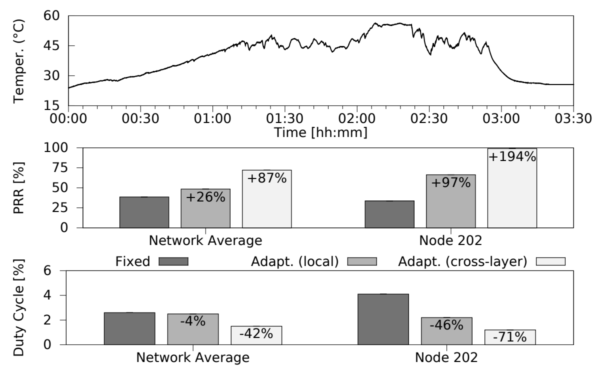


Fig. 10. Regeneration of a real-world outdoor temperature trace and impact on PRR and duty cycle on a network level and on a single node.

We finally use TempLab to time-lapse a 24-hours trace recorded in an outdoor deployment [4], and see what is the impact in a network with a density of one node every 8 $m^2$ when using $T^{'}_{CCA} = n^{'}_{f} + 6$ dBm. The results show that the adaptive approaches that we proposed significantly improve performance, both on a link basis and on a network level. Fig. 10 shows that the network sustains up to 42% lower radio duty cycle and 87% higher PRR in the presence of temperature variations commonly found in outdoor deployments, and that a single link may experience even up to 71% lower duty cycle and 194% higher packet reception rate.

## VI. Related Work

Several outdoor deployments and experimental studies have highlighted the impact of temperature on the quality of communications in wireless sensor networks. Bannister et al. [12] have reported that high temperatures can decrease the strength of the wireless signal. Wennerström et al. [4] have found experimental evidence of this problem on a long-term outdoor deployment. Boano et al. [3] have shown that the transmission power of communications at low temperatures can be safely decreased without deteriorating the performance of the network, and have precisely characterized the attenuation in received signal strength on different platforms [5]. All these works, however, simply report the degradation of the wireless signal as a consequence of an increase in temperature and do not provide a deeper analysis of what the implications are on communication protocols when operating a network outdoors.

Keppitiyagama et al. [22] have presented a poster showing that network protocols are affected by temperature and proposed to enhance them with temperature hints. In our earlier work, we have presented TempLab, a testbed infrastructure to study the impact of temperature on communication protocols [6], and used it to confirm the low performance at high temperatures. In this work, we exploit this testbed infrastructure to analyse *why* communication protocols are affected and propose *how* to mitigate the problems by dynamically adapting the CCA threshold to temperature variations.

Several works have suggested the use of adaptive CCA thresholds in the context of interference mitigation. Bertocco et al. [23] have provided hints for an optimal threshold selection in the presence of in-channel additive white Gaussian noise interference. Yuan et al. [24] have proposed to adjust the CCA threshold in the presence of heavy interference to reduce the amount of discarded packets due to channel access failures. Xu et al. [25] have designed a mechanism that dynamically adjusts the CCA threshold to enable concurrent transmissions on adjacent non-orthogonal channels and achieve high throughput.

Sha et al. [26] have studied the effects of the CCA threshold setting in noisy environments, and shown that interference can increase the number of false wake-ups in low-power-listening MAC protocols. To remedy this problem, they have proposed AEDP, an adaptive protocol that adjusts the CCA threshold in response to changes of ETX. While we share the idea that the CCA threshold cannot be set to an arbitrary value at compile-time, there are considerable differences with our work. First, AEDP is designed to achieve a desired performance in noisy environments and does not take into account the role of temperature. This may lead to problems, as AEDP requires an estimate of the noise floor and of the average RSSI value of all incoming links, which may change as temperature changes. Second, AEDP does not require a temperature model to adapt the CCA threshold, but instead requires information of observed interference in recent packet transmission attempts. In event-based networks, the reliance on ETX values may be a problem since packet transmissions are sparse.

An alternative approach to mitigate the impact of temperature may consist in increasing the transmission power at high temperatures, as suggested by the data-sheets of some radio chips. Although this would lead to an increased energy-consumption, it may simply not be possible: a node could already be using its highest power level. Furthermore, increasing the power based on the local temperature would only make the transmitted signal stronger, but would not solve the attenuation on the receiver side. Hence, our approach based on the signal strength attenuation modelling is more generic and effective.

## VII. Conclusions

The central tenet of our study is that temperature variations affect the efficiency of clear channel assessment and may compromise the operations of data link layer protocols based on carrier sense. We have shown that a reduced effectiveness of CCA at high temperatures compromises the ability of a node to avoid collisions and to successfully wake up from low-power mode. We have designed and evaluated two mechanisms to mitigate the problem by dynamically adapting the CCA threshold to temperature changes: one based on the temperature measured locally, and one on the highest temperature measured across all neighbouring nodes. Through an extensive experimental evaluation, we have shown that the proposed approaches increase the robustness of existing protocols to temperature variations and significantly improve the performance both on a link basis and on a network level.

## References


[1] R. Szewczyk et al., "Lessons from a sensor network expedition," *Wireless Sensor Networks*, vol. 2920, 2004.

[2] C. A. Boano et al., "Low-power radio communication in industrial outdoor deployments: The impact of weather conditions and ATEX-compliance," in *Proc. of the* $1^{st}$ *SensAppeal*, 2009.

[3] C. A. Boano, J. Brown, N. Tsiftes, U. Roedig, and T. Voigt, "The impact of temperature on outdoor industrial sensornet applications," *IEEE Transactions on Industrial Informatics*, vol. 6, no. 3, 2010.

[4] H. Wennerström, F. Hermans, O. Rensfelt, C. Rohner, and L.-A. Nordén, "A long-term study of correlations between meteorological conditions and 802.15.4 link performance," in *Proc. of the* $10^{th}$ *SECON*, 2013.

[5] C. A. Boano et al., "Hot Packets: A systematic evaluation of the effect of temperature on low power wireless transceivers," in *Proc. of the* $5^{th}$ *ExtremeCom*, 2013.

[6] C. A. Boano, M. Zúñiga, J. Brown, U. Roedig, C. Keppitiyagama, and K. Römer, "TempLab: A testbed infrastructure to study the impact of temperature on WSNs," in *Proc. of the* $13^{th}$ *IPSN*, 2014.

[7] J. Beutel, B. Buchli, F. Ferrari, M. Keller, L. Thiele, and M. Zimmerling, "X-SENSE: Sensing in extreme environments," in *Proc. of DATE*, 2011.

[8] C. Park, K. Lahiri, and A. Raghunathan, "Battery discharge characteristics of wireless sensor nodes: An experimental analysis," in *Proc. of the* $2^{nd}$ *SECON*, 2005.

[9] W. Guo, W. M. Healy, and M. Zhou, "Experimental study of the thermal impacts on wireless sensor batteries," in *Proc. of the* $10^{th}$ *ICNSC*, 2013.

[10] Q. Li, M. Martins, O. Gnawali, and R. Fonseca, "On the effectiveness of energy metering on every node," in *Proc. of the* $10^{th}$ *DCOSS*, 2013.

[11] L. Wanner, C. Apte, R. Balani, P. Gupta, and M. Srivastava, "Hardware variability-aware duty cycling for embedded sensors," *IEEE Transactions on VLSI Systems*, vol. 21, no. 6, 2013.

[12] K. Bannister, G. Giorgetti, and S. K. Gupta, "Wireless sensor networking for hot applications: Effects of temperature on signal strength, data collection and localization," in *Proc. of the* $5^{th}$ *HotEmNets*, 2008.

[13] C. A. Boano, T. Voigt, N. Tsiftes, L. Mottola, K. Römer, and M. Zúñiga, "Making sensornet MAC protocols robust against interference," in *Proc. of the* $7^{th}$ *EWSN*, 2010.

[14] J. Polastre, J. Hill, and D. E. Culler, "Versatile low-power media access for wireless sensor networks," in *Proc. of the* $2^{nd}$ *SenSys*, 2004.

[15] M. Buettner, G. V. Yee, E. Anderson, and R. Han, "X-MAC: A short preamble MAC protocol for duty-cycled wireless sensor networks," in *Proc. of the* $4^{th}$ *SenSys*, 2006.

[16] A. Dunkels, "The contikimac radio duty cycling protocol," Swedish Institute of Computer Science, Tech. Rep. T2011:13, 2011.

[17] D. Moss and P. Levis, "BoX-MACs: Exploiting physical and link layer boundaries in low-power networking," Stanford University, Tech. Rep. SING-08-00, 2008.

[18] C. A. Boano, T. Voigt, C. Noda, K. Römer, and M. Zúñiga, "JamLab: Augmenting sensornet testbeds with realistic and controlled interference generation," in *Proc. of the* $10^{th}$ *IPSN*, 2011.

[19] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *Proc. of the* $4^{th}$ *IPSN*, 2005.

[20] T. Winter et al., "RPL: IPv6 routing protocol for low-power and lossy networks," IETF, Tech. Rep., 2012.

[21] P. Levis, T. Clausen, J. Hui, O. Gnawali, and J. Ko, "The trickle algorithm," IETF, Tech. Rep., 2011.

[22] C. Keppitiyagama, N. Tsiftes, C. A. Boano, and T. Voigt, "Temperature hints for sensornet routing," in *Proc. of the* $11^{th}$ *SenSys*, 2013.

[23] M. Bertocco, G. Gamba, and A. Sona, "Experimental optimization of CCA thresholds in wireless sensor networks in the presence of interference," in *Proc. of the IEEE EMC*, 2007.

[24] W. Yuan, J.-P. M. Linnartz, and I. G. M. M. Niemegeers, "Adaptive CCA for IEEE 802.15.4 wireless sensor networks to mitigate interference," in *Proc. of IEEE WCNC*, 2010.

[25] X. Xu, J. Luo, and Q. Zhang, "Design of non-orthogonal multi-channel sensor networks," in *Proc. of the* $30^{th}$ *ICDCS*, 2010.

[26] M. Sha, G. Hackmann, and C. Lu, "Energy-efficient low power listening for wireless sensor networks in noisy environments," in *Proc. of the* $12^{th}$ *IPSN*, 2013.