Experimental Comparison of Frequency Hopping Techniques for 802.15.4-based Sensor Networks

Luca Stabellini and Mohammad Mohsen Parhizkar
The Royal Institute of Technology (KTH), ICT/Communication Systems
Electrum 418, SE-164 40 Kista, Sweden
Email: {lucast,mmpa}@kth.se

***Abstract*—Frequency hopping communication schemes represent an attractive alternative for interconnecting low power wireless sensor nodes operating in unlicensed bands. The use of multiple communication channels can in fact mitigate the negative effects of interference induced by collocated wireless networks and potentially results in improved reliability. With this respect, quite a few adaptive variations, aiming at improving the resilience of frequency hopping toward interference have been recently proposed. In this paper we present the experimental evaluation of three different hopping schemes: we implement a traditional hopping algorithm and two adaptive variations on TMote Sky sensor motes and quantify their energy performance under different channel conditions. We also compare the effectiveness of these three hopping techniques against the one of a communication scheme making use of a single channel. Our results, obtained considering a two-node topology, indicate that our previously proposed utility based adaptive frequency hopping approach is the most effective in avoiding interference and can significantly reduce the overall energy consumption despite its higher complexity. The performed experiments also show that even though reliable single-channel communication might be possible, by using frequency hopping sensors can limit the performance degradation induced by interference while avoiding the energy overhead introduced by the spectrum sensing algorithms that nodes have to use for identifying clear channels.**

***Keywords*-Frequency Hopping; Adaptive Frequency Hopping; Interference Mitigation; Coexistence in Unlicensed Bands; Wireless Sensor Networks;**

I. INTRODUCTION

A. Background

Frequency hopping communication techniques represent a common solution for interconnecting wireless personal area network devices operating in unlicensed bands. The basic idea implemented by these schemes is to allow communication among two or more wireless terminals by means of synchronous hopping over a defined set of channels (also referred to as the *hopset*) that are selected for packet transmissions in a pseudo-random fashion. Such a strategy guarantees a certain degree of frequency diversity and potentially allows to mitigate the interference that might be induced by transmissions of collocated wireless networks, consequently improving reliability.

This nice feature is extremely attractive for low-power devices such as wireless sensor nodes. As outlined by recent surveys conducted in the context of industrial automation (see for instance [1, 2]) the potential unreliability of wireless communications is in fact perceived as one of the major barriers to the adoption of wireless sensing technologies for commercial applications. By exploiting multiple communication channels through frequency hopping, sensor devices can mitigate the negative effects of interference and potentially improve communication reliability. We remark that the attention towards frequency hopping transmission schemes has been constantly growing during the last years as witnessed by the recent proliferation of radio standards and communication protocols adopting this solution: examples are provided by IEEE 802.15.1 [3], WirelessHART [4] and ISA SP100 [5].

While frequency hopping techniques can guarantee a certain resilience against bad channel conditions, it is well known that performance of this kind of systems can be severely degraded if some of the channels belonging to the used hopset constantly experience bad communication quality. For dealing with this problem adaptive algorithms have been proposed: in particular two approaches have been investigated in the literature. The first one (see for instance the adaptive specifications included in [6], [7] and references therein) aims at identifying bad channels that are subsequently removed from the used hopset whose cardinality is thus reduced: note that this approach is implemented by a variety of adaptive algorithms such as ISOAFH [8] (that targets the identification of interference induced by WLAN devices) and EAFH [9] (that also adapts the size of transmitted packets to the particular channel conditions of each frequency band). The latter instead adopts a probabilistic approach that rather than removing channels from the hopset, uses all the available frequency bands but with probabilities that depend on channel quality (see for instance [10] and [11]).

B. Problem Formulation and Contribution

These two approaches introduce different overhead, present different complexity and provide different advantages. For instance removing *bad* channels from the hopset results in relatively low complexity: on the other hand this choice might introduce delays in the adaptation process (due to the need of identifying those bad channels with a certain accuracy) and frequency bands that are removed from the hopset might have to be periodically re-checked resulting in additional wastes of energy and time. The probabilistic approach introduced in [10] allows to overcome these limitations: adaptation can in fact start immediately after the first packets are transmitted/received and the available resources can be exploited in a more granular manner. This however results in higher computational complexity and requires frequent exchange of information among the communicating nodes in order to maintain synchronous hopping and avoid the multi-channel hidden terminal problem [12].

We remark that the impact of these different design choices over the performance of wireless devices has always been evaluated through simulations, and we are not aware of any published research work aiming at quantifying and comparing through experiments on real hardware the effectiveness of different hopping algorithms. In this paper we provide such an experimental comparison. In particular, our contribution is two-fold:

- first, using TMote Sky sensor nodes we implement the two adaptive approaches previously described as well as a traditional hopping algorithm and quantify and compare their energy performance by means of extensive experiments under different channel conditions;
- we further compare the energy performance of frequency hopping against the one of a communication scheme making use of a single channel.

The use of (adaptive) frequency hopping has been envisaged for improving the performance of wireless systems under three different settings: (i) in presence of frequency static interference (such as for instance the one induced by IEEE 802.11 b/g devices), (ii) in presence of frequency dynamic interference (such as the one induced by collocated networks making use of frequency hopping) and (iii) in presence of bad channel conditions (for instance induced by multipath fading, frequency selective channel responses or other propagation anomalies). In this work we consider only frequency static interference. As done in many other papers (see for instance [13, 14]) we limit the focus of our investigation to a simple two-node topology: the extension to networks comprising more than two nodes is left for future work.

The rest of this paper is organized as follows. Section II outlines the setting of our experiments and Section III describes the hopping algorithms we implemented. Experimental results are presented in Section IV, while conclusions are drawn in Section V.

II. Experimental Setup

A. Network Scenario

The setting of our experiments is sketched in Figure 1: we focus on a simple two-node topology comprising two TMote Sky sensor nodes S_1 and S_2 located 1 meter far away from each other, and consider the exchange of a certain bulk of data, consisting of N packets, from node S_1 (acting as transmitter) to node S_2 (the receiver). Each transmitted packet has a payload of 100 Bytes. The two sensors run the Contiki operating system [15] and are connected to two PCs through a USB connection that allows to collect transmission statistics. Packet transmissions are implemented using the following handshake mechanism: on slot i, node S_1 sends a data block; S_2 verifies the correctness of the received packet by means of a 16-bit cyclic redundancy check (we used the CRC16 provided by the Contiki operating system [15]). An acknowledgement or a not acknowledgement (requesting the retransmission of the corrupted block) is then transmitted on slot $i+1$.

The TMote Sky used for our experiments feature an IEEE 802.15.4 2420 Chipcon wireless transceiver operating in the 2.4 GHz ISM band: the available hopset comprises thus the 16 frequency bands $c_{11}, \ldots, c_{26}$ specified by the IEEE 802.15.4 radio standard. We implemented a simple MAC-layer synchronization routine where nodes hop to the channel that has to be used for the upcoming transmission either immediately after sending or receiving a packet or after a certain time-out t_{Max} expires: for our experiments we fixed $t_{\text{Max}} = 25$ [ms].

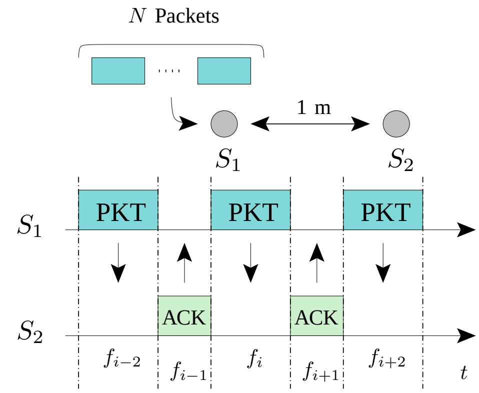

Fig. 1. Sketch of the considered scenario.

We focused on a simple two-node topology for two reasons. On one hand, the implementation of frequency hopping schemes over multi-node and potentially multi-hop networks requires that different issues, one of them being synchronization, are addressed. This is out of the scope of this paper: furthermore, we note that the same problems will arise independently on the used hopping technique. Considering only two sensors simplifies the implementation process and allows us to focus on the comparison of the energy performance of the different communication techniques. As a second aspect, we remark that networks comprising several sensors can potentially be organized in a countless number of different topologies. The choice of a particular topology (for instance a star rather than a tree or a mesh) might make the obtained results dependent on the particular considered setting: by focusing on the single link between two nodes instead it is possible to obtain general results that are not topology dependent.

B. Experimental Approach

We performed two different experimental campaigns. For the first one, we selected an interference-free environment: we set the transmission power of the nodes so as to achieve a negligible packet loss rate (we verified that an output power of -10 dBm was sufficient for this purpose) and we *artificially* controlled using software-defined values the probability p_i of receiving a corrupted packet over channel c_i (note that in fact all packets are correctly received however with probability p_i, a packet is discharged and considered lost). This approach, that has previously been used for instance in [16, 17], basically permits to *simulate* the performance of the considered hopping algorithms on real motes (thus allowing to quantify their exact energy consumption) while controlling the packet error probability experienced over the wireless channel.

Our second campaign was instead performed inside the office spaces of the Radio Communication Systems department of KTH where the 2.4 GHz ISM band is heavily used by several wireless terminals such as laptop, PDA and wireless keyboards/mouses: as an example, the variation of average channel occupancy over the 16 IEEE 802.15.4 channels during a 7-day period is shown in Figure 2. The spectrum is mainly utilized by WLAN devices (i.e., operating within the IEEE 802.11g radio standard): on the plotted figure three non-overlapping WiFi carriers can be easily identified.

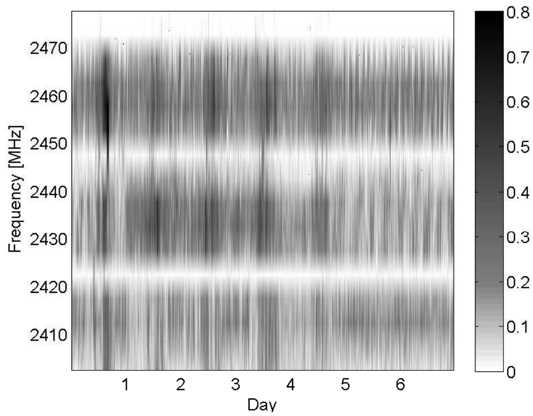

Fig. 2. Average channel occupancy for the 16 IEEE 802.15.4 channels during a 7-day period, from April 19 to April 25.

Transmissions of these devices, although not controllable, provide an example of interference pattern that sensors are likely to experience in real scenarios and thus represent an excellent source of interference for our motes. Our experiments have been performed over a time-frame of seven days: on each day we iterated several transmissions for each of the considered hopping algorithms in order to ensure that all of them were tested under a wide range of channel conditions. The purpose of this second campaign was to *qualitatively* assess the performance of the hopping techniques in a real scenario. Our conclusions are however mainly based on data obtained during the first round of controlled experiments.

C. Performance Metrics

We quantified the performance of the considered communication schemes by measuring the total energy E^{Tot} spent by the two-node system for the successful delivery of the specified bulk of data: this accounts for the energy spent while transmitting and receiving packets and control messages as well as for the energy required by the CPU of the two nodes. For this purpose, we used the online energy estimation routine [18] provided by the Contiki operating system: this allows to measure the time spent by nodes on each of the four following states: transmit, receive, CPU and LPM (Low Power Mode). The total consumed energy can then be computed multiplying the obtained times by the power consumption of sensors on each state (for TMote Sky we have: listen 60 [mW], transmit (-10dBm) 33 [mW], CPU 5.4 [mW] and LPM 0.1635 [mW] [19]). We normalized the obtained values to the amount of energy required to complete a packet exchange (comprising both the transmission of the data packet as well as the following acknowledgement) in interference-free conditions.

III. Hopping Algorithms

We now briefly describe the three hopping algorithms that have been the object of our evaluation. In particular we implemented:

- a traditional Frequency Hopping (FH) scheme;
- an Adaptive Frequency Hopping algorithm similar to the one defined in [6]; this adaptive approach is the one adopted by several radio standards such as for instance WirelessHART [4] and IEEE 802.15.1 [3];
- the Utility Based Adaptive Frequency Hopping (UBAFH) algorithm introduced in [10].

We also considered a simple communication scheme where only a single channel is used and no hopping strategy is implemented. The aforementioned hopping approaches will be detailed in the next subsections.

A. Traditional Frequency Hopping

If a traditional frequency hopping technique is implemented the channels belonging to the hopset are used in a pseudo-random fashion. For this purpose S_1 and S_2 share a common seed: this is used to generate random numbers and chose the frequency band that shall be used for the upcoming transmission. All the 16 available channels are equally likely to be selected in each time-slot.

B. Adaptive Frequency Hopping

We consider the adaptive frequency hopping algorithm specified in [6] (note that in [6] the focus was on the IEEE 802.15.1 radio standard: we here generalize the proposed scheme to IEEE 802.15.4): S_1 and S_2 estimate the packet error rate experienced on each channel using a certain number of transmissions N_E. In this way S_1 estimates the probability of receiving a corrupted ACK/NACK, while S_2 estimates the probability of receiving a corrupted data packet. After this *channel classification* procedure has been completed, S_1 reports to S_2 his estimates, S_2 computes average channel conditions (by averaging his estimates with the ones received from S_1) and updates the hopset by removing channels with packet error rate greater than a certain threshold p^{Max}. The updated hopset is then communicated to S_1 and adaptation can start. This procedure can eventually be repeated on a periodic fashion in order to deal with changes of channel conditions.

We remark that [6] do not specifies the values of N_E and p^{Max} which can therefore be vendor specific: for our experiments we assumed $N_E = 16 \cdot 20$ (thus channel conditions are estimated considering in average 20 transmissions for each of the available frequency bands) and $p^{\text{Max}} = 0.5$. We stress that different choices for these parameters can be used to implement different performance tradeoffs. A low value of N_E allows to shorten the time required to perform channel classification and thus reduces the adaptation delay: on the other hand if N_E is too small, channels might be classified in an inaccurate manner and for instance good frequency bands might erroneously be removed from the hopset while bad channels might not be properly identified. Similar considerations should be made when selecting the packet error rate threshold p^{Max}: a high threshold might lead nodes to hop over interfered frequencies while lower values might induce a very selective channel classification procedure where several channels are removed from the hopset decreasing the degree of frequency diversity. This might be undesirable if nodes experience both frequency static interference as well as multi-path fading. The value we assumed for our experiments i.e. $p^{\text{Max}} = 0.5$ has been suggested in [13] and has been used in other published works.

C. Utility Based Adaptive Frequency Hopping

The utility based adaptive frequency hopping algorithm proposed in [10] adopts a different approach: S_1 and S_2 constantly maintain estimates $\hat{p}(c_i)$ for the packet error rate experienced on each of the available frequency bands. These estimates are computed using a window moving average that evaluates $\hat{p}(c_i)$ over channel c_i accounting for the last $N_T = 32$ transmissions. The obtained values are then mapped to a probability mass function defining channel usage probabilities and assigning to channels with better conditions higher values. For complexity reasons we modified the mapping function defined in [10] and considered instead:

$$f : \hat{p}(c_i) \rightarrow f\left(\hat{p}(c_i)\right) = \frac{\nu(\hat{p}(c_i))}{\sum_{j=1}^{16} \nu(\hat{p}(c_j))} \quad (1)$$

where:

$$\nu(\hat{p}(c_i)) = \begin{cases} 20 \cdot (1 - \hat{p}(c_i)) \cdot 32 & \text{if } \hat{p}(c_i) \leq \frac{3}{32} \\ 5 \cdot (1 - \hat{p}(c_i)) \cdot 32 & \text{if } \frac{3}{32} < \hat{p}(c_i) \leq \frac{12}{32} \\ 3 & \text{if } \hat{p}(c_i) > \frac{12}{32} \end{cases} \quad (2)$$

Note that the factor 32 that multiplies $1 - \hat{p}(c_i)$ is introduced in order to obtain integer quantities and reduce computational complexity. To the payload of each packet S_1 and S_2 add two bytes (thus the payload in this case has a total size of 102 octets) containing the outcomes of the last 16 packets transmissions: this allows to keep synchronous estimates of packet error rate at the two nodes (the reader is referred to [10] for additional details). The channel to be used at time-slot k is then selected using the information that nodes have up to time-slot $k - 16$. Also in this case, synchronous channel selection is ensured by using a seed known to both nodes. Note that channels are still chosen in a pseudo-random fashion, however, while for a traditional hopping algorithm, all the channels are equally likely to be used, in this case channels with better conditions (i.e. lower packet error rate) are assigned higher usage probabilities (proportionally to $\nu(\hat{p}(c))$) and are consequently selected more often than frequency bands where nodes experience high packet error rate.

IV. Results

A. Interference Controlled Environment

We start our performance evaluation by quantifying the complexity added by the adaptive schemes in absence of interference. Under these conditions, the energy consumption of the traditional FH algorithm represents our reference case: in Figure 3 (top-left) we show the relative amount of energy consumed by the two-node system while in the CPU, transmitting and receiving states. Note that energy spent while receiving represents the major component. This is due to the fact that receiving (or idle listening) is more energy costly than transmitting; furthermore, prior to each packet transmission, the data that are to be sent have to be copied from the micro-controller to the radio transceiver: during this operation the radio of sensors is in the listening state, and as a result the time spent while listening is greater than the one spent transmitting (see [20] for additional details). On the top-right side of Figure 3 we consider instead the utility based adaptive frequency hopping algorithm proposed in [10]: the total energy consumption is in this case increased by approximatively 4%. This is due both to longer listening and transmitting times (note that nodes add to each transmitted packet a two-byte field for synchronization purposes) as well as to the higher computational complexity of the adaptive procedure which results in increased CPU energy consumption. The adaptive frequency hopping algorithm described in Section III-B basically presents the same energy performance as traditional FH (since the adaptation procedure we implemented is on demand, and in absence of interference no adaptation is performed, see the bottom-left plot in Figure 3). Finally, if the single channel approach is selected, the overall energy consumption is reduced by approximatively 7%: this is due to lower complexity (no generation of random numbers is performed) as well as to the fact that nodes do not need to switch frequency band after transmitting/receiving packets and acknowledgements[1].

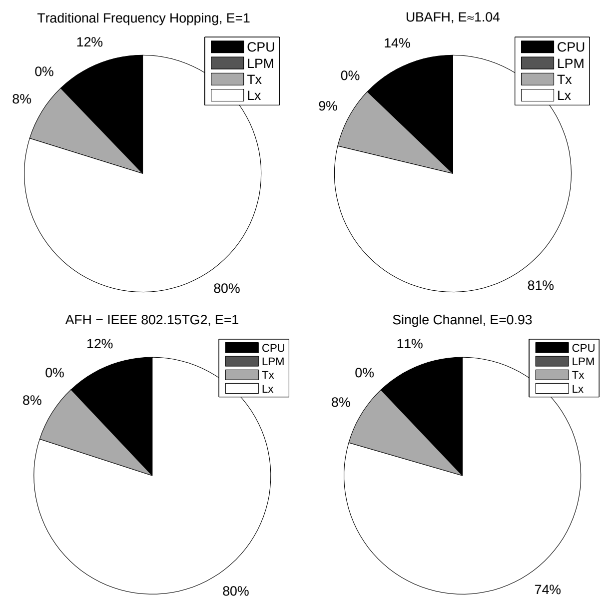

Fig. 3. Relative energy consumption in the four different energy states of the two node system for FH, UBAFH, AFH and Single Channel scheme in interference free conditions. Note that percentages of UBAFH and of the Single Channel scheme sum up to 104% and 93% respectively since we normalized the obtained values to the energy consumption of FH.

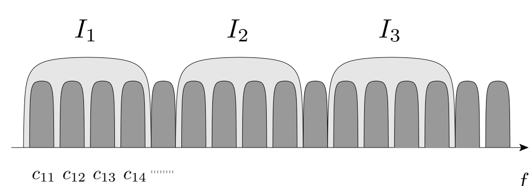

Fig. 4. Interference Scenario. Each interfering carrier I_i induces a certain packet error probability over the overlapping IEEE 802.15.4 channels.

Let us now start our performance evaluation. Using the methodology described in Section II we emulated the presence of three WLAN carriers (I_1, I_2, I_3, see Figure 4) overlapping

[1]For the CC2420 radio unit, channel switching time is in the order of 200 μs [19] and it is equivalent to the time required to transmit about 50 bits

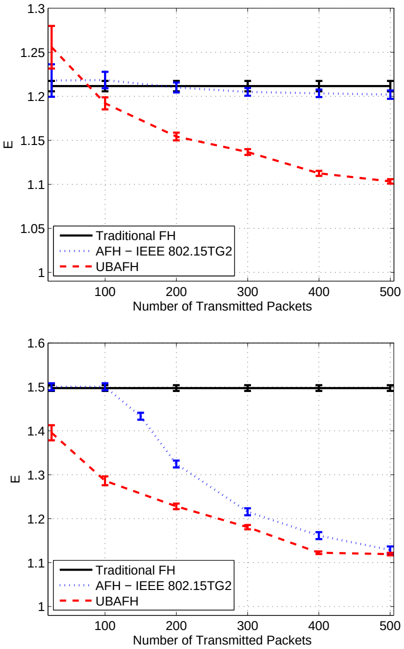

Fig. 5. Results for Scenario 1. Average energy per packet for $p = 0.4$ (top) and $p = 0.8$ (bottom).

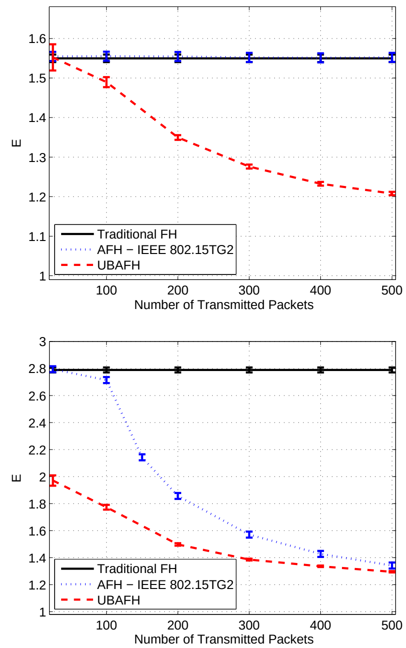

Fig. 6. Results for Scenario 2. Average energy per packet for $p = 0.4$ (top) and $p = 0.8$ (bottom).

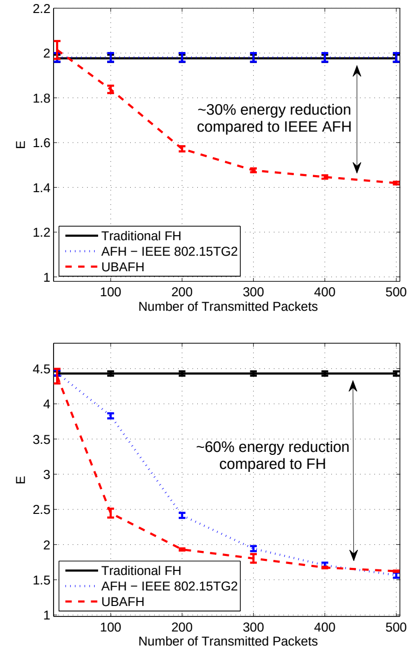

Fig. 7. Results for Scenario 3. Average energy per packet for $p = 0.4$ (top) and $p = 0.8$ (bottom).

with the channels used by sensors and we varied the packet error rate experienced over those channels. In particular we considered the following scenarios:

1) Scenario 1: only one WLAN carrier (I_1) is active. This overlaps with the IEEE 802.15.4 channels $11 - 14$;
2) Scenario 2: two WLAN carriers (thus both I_1 and I_2) are active. I_1 overlaps with channels $11 - 14$, while I_2 overlaps with channels $16 - 19$.
3) Scenario 3: all three WLAN carriers are active. These overlap respectively with channels $11 - 14$, $16 - 19$ and $21 - 24$.

For each of these scenarios we run our experiments for two different settings. In the first one, the packet error probability induced by the WLAN carriers is set to $p = 0.4$ while for the latter we consider $p = 0.8$: these two values are respectively below and above the threshold packet error probability used by the channel classification procedure defined by IEEE AFH (see Section III-B). In all cases, we considered symmetric channel conditions at the two nodes i.e. nodes experience equal packet error probabilities on the same channel.

Results for Scenarios 1, 2 and 3 are respectively presented in Figures 5, 6 and 7, where we show as a function of the amount of transmitted data N the average energy per packet E for the three hopping schemes. 95% confidence intervals are also plotted in all curves. While the energy performance of traditional frequency hopping do not significantly depend on the amount of transmitted data, the other algorithms can benefit from adaptation and in fact transmitting a larger amount of packets allows to improve energy efficiency. It should be remarked how the different adaptive approaches implemented by the two schemes we considered lead to different energy performance. The traditional adaptive algorithm proposed in [6], makes use of an ineffective channel classification procedure that allows the adaptation process to start only after a significant number of packets has been transmitted. We further remark that the use of a *binary* approach, where channels are either used for hopping or completely removed from the hopset, is very sensitive to the choice of the used threshold. Over a set of channels presenting only frequency static interference and for an appropriate packet error rate threshold, this adaptive strategy provides the best performance since nodes hop only over clear frequency bands: however energy efficiency can easily be deteriorated if the value of p^{Max} is not properly chosen. In our experiments, where we selected on purpose an improper threshold, a packet error probability equal to 0.4 was sporadically allowing to classify the considered frequency bands as interfered, preventing the algorithm from adapting. This behavior can potentially be improved by lowering the threshold used by the channel classification procedure, however the same problem might arise also with a lower value of p^{Max} if on some of the available channels nodes experience a packet error rate that is just slightly below the new threshold.

The probabilistic approach adopted by UBAFH overcomes these limitations. As shown by the plotted curves, adaptation can start as soon as a few packets are transmitted: this results in lower energy consumption. Moreover, the implemented algorithm allows a more granular exploitation of the available resources if compared to the binary strategy implemented by IEEE AFH. This is clearly shown by the energy performance in presence of low interfering activities (top plots of Figures 5, 6 and 7): channels experiencing low (but still significant) packet error rates are used less frequently than not-interfered channels and this allows to reduce energy

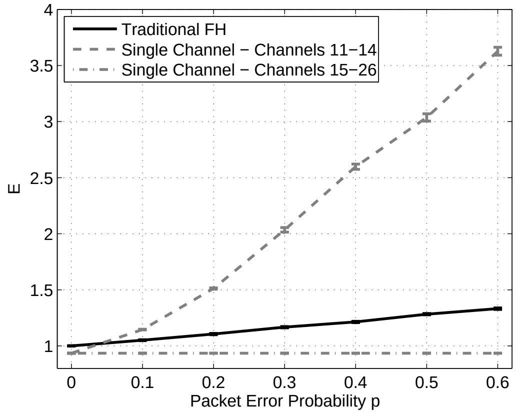

Fig. 8. Average Energy per packet as a function of the experienced packet error probability for frequency hopping and single channel scheme. Results are obtained considering $N = 100$ packets.

consumption of up to 30% if compared to IEEE AFH.

In order to compare the energy performance of frequency hopping against the one of a communication scheme making use of a single channel we performed a very simple experiment: we activated the WLAN carrier overlapping with the first four 802.15.4 channels (I_1 with reference to Figure 4) and varied the packet error probability experienced by sensors in the range $[0, 0.6]$. The average energy per packet for frequency hopping and single channel approach are shown in Figure 8. Note that if nodes operate over an interfered frequency, packet error probabilities as low as 10% are already sufficient to justify the use of frequency hopping proving that the overhead introduced by channel hopping is relatively small.

B. Real Environment

Results obtained in a real and uncontrolled wireless scenario validate the considerations made in the previous sub-section. Average energy per packet for FH, IEEE AFH and UBAFH are shown in Figure 9: the two plots are obtained considering the transmission of bulks of data consisting of 100 (top) and 500 (bottom) packets. For each algorithm we performed 200 experiments per day, 100 between 10 to 12 AM and 100 between 2 to 4 PM: during these hours the 2.4 GHz ISM band was mainly used by WiFi devices. For a relatively small amount of data, the three algorithms basically perform in the same way and lead to similar energy consumptions. However, while more and more packets are transmitted, adaptation plays an important role as shown in the bottom plot of Figure 9. IEEE AFH basically fails in identifying bad channels (in fact, only in a few cases we observed during the channel classification phase a packet error rate greater than the fixed threshold): these are consequently kept in the hopset and used as often as the good ones. The approach implemented by UBAFH instead allows to progressively decrease the probability of selecting frequency bands where sensors experience bad conditions and this results in lower energy consumption.

We finally compare always in the office spaces of the radio communication systems department of KTH the performance of the considered frequency hopping techniques against the one of a communication scheme making use of a single channel. Results for this scenario are presented in Figure 10 where we show the average energy per packet for the different communication schemes. For the single channel case, energy values are presented for all the 16 available frequency bands. For each of them, we performed 25 experiments per day (between 2 to 4 PM) and repeated these experiments on 7 different days. Note that on channels that overlap with the WiFi carriers used for internet access in the environment of our evaluation (see Figure 2), channel conditions can be extremely bad and energy consumption can be increased of up to 6 times. The use of frequency hopping allows to mitigate these problems by *averaging* channel conditions and reducing the high energy consumption that nodes experience in the worst case single-channel scenario.

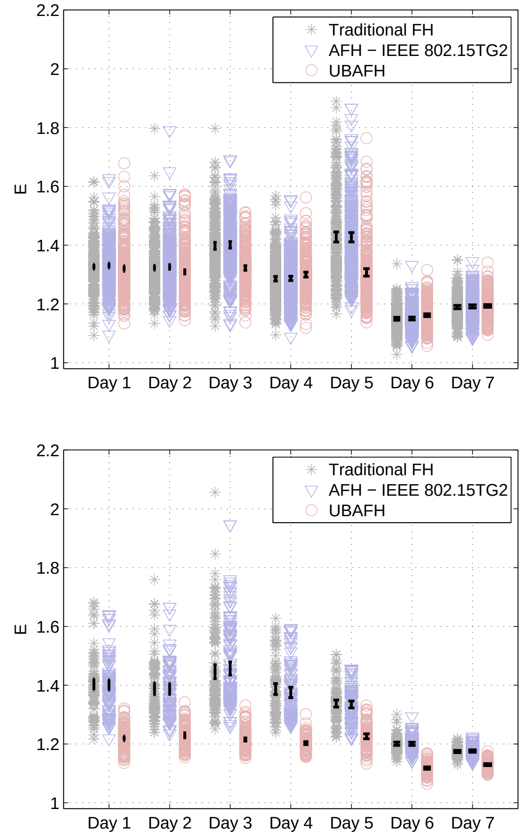

Fig. 9. Energy performance for traditional FH, IEEE AFH and UBAFH in a real environment. Results have been obtained considering 100 (top) and 500 (bottom) packet transmissions. 95% confidence intervals are also shown. Note that day 6 and 7 correspond to Saturday and Sunday

We stress the importance of this last observation: recently published works (see for instance [21]) have questioned the utility of frequency hopping schemes in real environments pointing out that in typical settings, when multiple channels are available, it is likely that there is a non-empty set of clear and not interfered frequency bands. We remark however that while identifying those channels by means of dedicated spectrum sensing algorithms might be quite straightforward [22], the energy overhead introduced by this procedure might be significant and could be equivalent to the energy required to

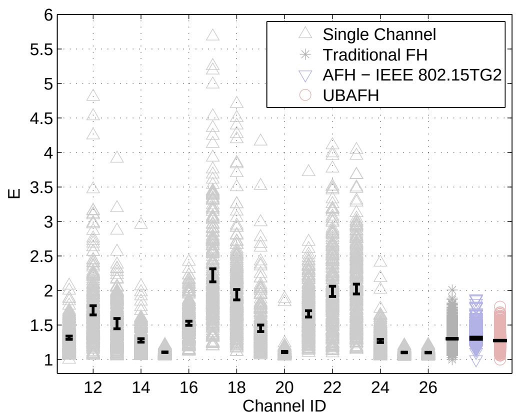

Fig. 10. Average Energy per packet for the single channel approach and for the three considered frequency hopping schemes. Results are obtained transmitting $N = 100$ packets. 95% confidence intervals are also shown.

transmit several tens of packets (see [17]). Thus, even though reliable single-channel communication is indeed possible (note for instance that in Figure 10, if channels 15, 20, 25 or 26 are selected, average energy per packet is lower than the one achieved by the frequency hopping scheme), the use of frequency hopping allows to limit the performance degradation that can potentially be induced by interference and at the same time permits to avoid the energy overhead introduced by the spectrum sensing algorithms that nodes have to run in order to identify clear channels.

V. Conclusions

In this paper we presented the experimental evaluation of different frequency hopping techniques for wireless sensor networks. Using TMote Sky sensor nodes operating within the IEEE 802.15.4 radio standard we implemented a frequency hopping algorithm and two adaptive schemes and quantified their complexity and energy performance under different channel settings. Our results have shown that traditional frequency hopping schemes, where channels are used in a pseudo-random fashion are very sensitive to interference that can severely degrade their energy efficiency. On the other hand, the utility based adaptive frequency hopping algorithm we recently proposed, introduces significant computational complexity but allows to effectively adapt the hopping pattern in presence of bad channel conditions. In our experiments this led to energy savings of up to 60 percent if compared to non-adaptive schemes and as high as 30 percent if compared to the other and more traditional adaptive approach that was considered in our evaluation. Comparison with a traditional communication scheme using a single channel has also outlined that frequency hopping is very useful in presence of interference and can be exploited in order to limit the performance degradation induced by transmissions of collocated wireless devices.

Promising directions for future work might include the extension of our experimental campaign for investigating the behavior of the considered hopping schemes in presence of frequency dynamic interference and propagation anomalies (such as the multi-path or frequency selective fading that might for instance arise in industrial settings). It could also be interesting to evaluate the effectiveness of the considered hopping techniques on networks comprising more than two nodes.

References

[1] J. Morse, “Market Pulse: Wireless in Industrial Systems: Cautious Enthusiasm“, Embedded Systems, Winter 2006.
[2] “WSN for Smart Industries: A Market Dynamics Report“, OnWorld, September 2007.
[3] “Part 15.1: Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs)“, ANSI/IEEE Standard 802.15.1-2005.
[4] “Why Wireless HART? The Right Standard at the Right Time“, white paper, October 2007. Available online at www.hartcomm2.org.
[5] “The ISA100 Standards - Overview & Status“, October 2008, Available online at http://www.isa.org.
[6] “Part 15.2: Coexistence of Wireless Personal Area Networks with other Wireless Devices Operating in Unlicensed Frequency Bands“, ANSI/IEEE Standard 802.15.2-2003.
[7] P. Popovski, H. Yomo, and R. Prasad, “Strategies for Adaptive Frequency Hopping in the Unlicensed Bands“ in IEEE Wireless Communications, Vol. 13, No. 6, December 2006.
[8] M. C.-H. Chek and Y.-K. Kwok, “Design and Evaluation of Practical Coexistence Management Schemes for Bluetooth and IEEE 802.11b Systems“, in Computer networks, Vol. 51, Issue 8, June 2007.
[9] A. C.-C. Hsu, D. S. L. Wei, C.-C. J. Kuo, N. Shiratori, and C-Ju Chang, “Enhanced Adaptive Frequency Hopping for Wireless Personal Area Networks in a Coexistence Environment“, in Proceeding of Global Telecommunications Conference (GLOBECOM), 2007.
[10] L. Stabellini, L. Shi, A. A. Rifai, J. Espino, and V. Magoula, “A New Probabilistic Approach for Adaptive Frequency Hopping“, in Proceedings of International Symposium on Personal Indor and Mobile Radio Communications, PIMRC, 2009.
[11] K. J. Park, T. R. Park, C. D. Schmitz, and L. Sha, “Entropy-Maximization Based Adaptive Frequency Hopping for Wireless Medical Telemetry Systems“, in Proceedings of the 1st ACM International Workshop on Medical-Grade Wireless Networks, 2009.
[12] J. So and N. Vaidya, “Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals Using A Single Transceiver“, in Proceedings of the Fifth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), 2004.
[13] N. Golmie, O. Rebala, and N. Chevrollier, “Bluetooth adaptive frequency hopping and scheduling“, in Proceedings of Military Communications Conference (MILCOM), Boston, USA, October 2003.
[14] K. J. Park, T. R. Park, C. D. Schmitz, and L. Sha, “Design of Robust Adaptive Frequency Hopping for Wireless Medical Telemetry Systems“, in IET Communications, Vol. 4, No. 2, 2010.
[15] A. Dunkels, B. Grönvall, and T. Voigt, “Contiki - a Lightweight and Flexible Operating System for Tiny Networked Sensors“, in Proceedings of the IEEE Workshop on Embedded Networked Sensors (Emnets-I), 2004.
[16] M. Rossi, G. Zanca, L. Stabellini, R. Crepaldi, A. F. Harris III, and M. Zorzi, “SYNAPSE: A Network Reprogramming Protocol for Wireless Sensor Networks using Fountain Codes“, in Proceedings of the IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2008.
[17] L. Stabellini and M. U. Javed, “Experimental Comparison of Dynamic Spectrum Access Techniques for Wireless Sensor Networks“, in Proceedings of Vehicular Technology Conference (VTC Spring), 2010.
[18] A. Dunkels, F. Österlind, N. Tsiftes, and Z. He, “Software-Based On-Line Energy Estimation for Sensor Nodes“, in Proceedings of the Fourth IEEE Workshop on Embedded Networked Sensors (Emnets IV), June 2007.
[19] ‘Tmote Sky Data Sheet‘ (2006) Moteiv, San Francisco, CA. Available online at: http://www.moteiv.com/products/docs/tmote-skydatasheet.pdf.
[20] F. Österlind and A. Dunkels, “Approaching the Maximum 802.15.4 Multi-Hop Throughput“, in Proceedings of the Fifth ACM Workshop on Embedded Networked Sensors (HotEmNets 2008), June 2008.
[21] J. Ortiz and D. Culler, “Multichannel Reliability Assessment in Real World WSNs“, in Proceedings of the 9th IEEE/ACM International Conference on Information Processing in Sensor Networks (IPSN), 2010.
[22] L. Stabellini and J. Zander, “Energy Efficient Detection of Intermittent Interference in Wireless Sensor Networks“, to appear in International Journal of Sensor Networks, 2010.