Adaptive Synchronization in IEEE802.15.4e Networks — Source link

David Stanislowski, Xavier Vilajosana, Qin Wang, Thomas Watteyne ...+1 more authors

Institutions: University of California, Berkeley

Related papers:

- OpenWSN: a standards-based low-power wireless development environment

- Tsmp: time synchronized mesh protocol

- Standardized Protocol Stack for the Internet of (Important) Things

- Adaptive synchronization in multi-hop TSCH networks

- Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH

Share this paper:

View more about this paper here: https://typeset.io/papers/adaptive-synchronization-in-ieee802-15-4e-networks-3wx4nmeefo

Document Version

This is the Accepted Manuscript version. The version in the Universitat Oberta de Catalunya institutional repository, O2 may differ from the final published version.

Enquiries

If you believe this document infringes copyright, please contact the Research Team at: repositori@uoc.edu

Adaptive Synchronization in IEEE802.15.4e Networks

David Stanislowski, Xavier Vilajosana, *Member, IEEE*, Qin Wang,
Thomas Watteyne, *Member, IEEE*, Kris Pister

Abstract—**Industrial low-power wireless mesh networks are shifting towards time synchronized medium access control (MAC) protocols which are able to yield over 99.9% end-to-end reliability, and radio duty cycles well below 1%. In these networks, motes use time slots to communicate, and neighbor motes maintain their clocks' alignment, typically within $1ms$. Temperature, supply voltage and fabrication differences cause the motes' clocks to drift with respect to one another. Neighbor motes need to re-synchronize periodically through pairwise communication. This period is typically determined *a priori*, based on the worst case drift.**

In this article, we propose a novel technique which measures and models the relative clock drift between neighbor motes, thereby reducing the effective drift rate. Instead of re-synchronizing at a preset rate, neighbor motes re-synchronize only when needed. This reduces the minimum achievable duty cycle of an idle network by a factor of 10, which, in turn, lowers the mote power consumption, and extends the network lifetime. This Adaptive Synchronization is implemented as part of IEEE802.15.4e in the OpenWSN protocol stack, and is validated through extensive experimentation.

Index Terms—**IEEE802.15.4e, Synchronization, Wireless Sensor Networks, Duty Cycle, Energy Consumption,TSCH.**

I. INTRODUCTION

In the last decade, contention-based wireless medium access control (MAC) layers have been used in many low-power wireless mesh protocols. ZigBee[1] has been the *de-facto* standard for these types of networks, but has failed to fulfill the industrial reliability requirements.

Things started to change with the development of the Time Synchronized Channel Hopping (TSCH) technique that was adopted by major industrial low-power wireless standards such as WirelessHART [1] and the recently as a part of the published IEEE802.15.4e standard [2]. As of today, several commercial networking providers are offering 99.9% reliable MAC layers that provide radio duty cycles well below 1%, thereby reducing the mote power consumption and increasing the network lifetime[2].

D. Stanislowski and K. Pister are with BSAC, University of California, Berkeley, CA, USA.

X. Vilajosana is with BSAC, University of California, Berkeley, CA, USA and Universitat Oberta de Catalunya, Barcelona, Spain and Worldsensing, Barcelona, Spain.

Q. Wang is with with BSAC, University of California, Berkeley, CA, USA and University of Science and Technology, Beijing, China.
T. Watteyne is with with BSAC, University of California, Berkeley, CA, USA and Dust Networks/Linear Technology, Hayward, CA, USA.

[1] http://www.zigbee.org/.
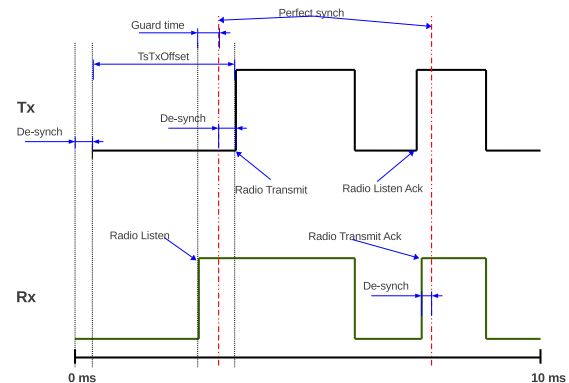
[2] http://www.linear.com/products/smartmesh_ip.

Fig. 1.   Timeline of an IEEE802.15.4e slot showing how two motes synchronize by exchanging a data packet.

The IEEE802.15.4e standard [2] is an amendment to the MAC protocol of IEEE802.15.4-2006 [3]. It achieves better reliability and lower power consumption through Time Synchronized Channel Hopping (TSCH). All motes in an IEEE802.15.4e network are synchronized and time is split into time slots, each typically $10ms$ long. Time slots are grouped into a super-frame which continuously repeats over time.

A schedule instructs each mote what to do in each time slot: send to a particular neighbor, receive from a particular neighbor, or sleep [4]. Channel diversity is obtained by specifying, for each send and receive slot, a channel offset. The same channel offset translates into a different frequency on which to communicate at each iteration of the super-frame. The resulting channel hopping communication reduces the impact of external interference and multi-path fading, thereby increasing the reliability of the network [5].

Crystal oscillators are commonly used for timing since they offer a good trade-off between power consumption, frequency stability and cost. The frequency of a crystal is affected by manufacturing differences, temperature, and supply voltage. A crystal oscillator is rated by its frequency stability: a $32kHz$ crystal rated $30ppm$ (parts-per-million) will pulse somewhere between $32768.99Hz$ and $32767.01Hz$. Two motes equipped with these crystal can *drift* by $60ppm$ to one another (one going fast, one going slow), i.e. they will desynchronize by $60\mu s$ each second. Motes therefore need to re-synchronize periodically.

In a TSCH network, all transmitting motes are scheduled to begin transmission at the same time in a slot (typically about

2ms into the slot, called the TsTxOffset). To allow for some de-synchronization, receivers start listening some time before this instant (see Fig. 1), and keep listening some time after. This duration is called the "guard time"

Assuming a guard time of $1ms$, and if two motes are equipped with $30ppm$ crystals, it takes $16s$ for them to desynchronize by more than $1ms$. Since they cannot communicate if they get desynchronized past the guard time, they periodically re-synchronize. Eq. (1) can be used to determine the maximum allowable synchronization period τ_{ka} as a function of the guard time T_g and the drift rate $\Delta\nu$. Re-synchronizing more frequently than the limit obtained with Eq. (1) guarantees that the motes stay synchronized within their guard time. The larger the guard time T_g, or the smaller the drift $\Delta\nu$, the less frequent motes need to re-synchronize.

$$\tau_{ka} = \frac{T_g}{\Delta\nu} \qquad (1)$$

As detailed in Section III, re-synchronizing in a TSCH network involves exchanging a data packet and an acknowledgment. Since the radio consumes power, it is best not to have to re-synchronize too often. We call the idle duty cycle the portion of time the radio of a mote needs to be on just to keep synchronized to its neighbors. The smaller this value, the longer the lifetime of the mote.

Typically, the worst-case drift rate is used to hard-code the re-synchronization period in the motes. This often results in "over-synchronization", hence a waste of energy. This article proposes a technique which allows a mote to measure its clock drift rate with a neighbor. Using this information, it tracks its neighbor's drift by periodically applying internal correction to its clock, which allows it to re-synchronizes less often. This estimated correction helps to ensure that even in very large networks the total de-synchronization between two arbitrary nodes is reduced as individual parent-child drifts are reduced.

This technique, called *Adaptive Synchronization*, is added to the TSCH mode of the IEEE802.15.4e implementation of the OpenWSN protocol stack [6], and used on real hardware. Experimental results presented in Section V show a reduction of the minimum achievable duty cycle of an idle network by a factor of 10.

II. RELATED WORK

Network synchronization is a well-studied topic. Lindsey *et. al.* published early theoretical work on synchronization, covering independent clocks settings, master-slave hierarchical settings (centralized), time reference distribution and mutual synchronization (decentralized) [7], [8]. The emergence of cellular and wireless networks came with new requirements on network synchronization, such as multi-frame synchronization [9] or network-wide synchronization in ad-hoc [10] networks. Recently, synchronization has also been studied in the context of industrial real-time systems, which introduces strict accuracy constraints [11]. The Time Synchronized Mesh Protocol (TSMP) [12] and the Timing-sync Protocol for Sensor Networks (TPSN) [13] were the first to apply synchronization techniques to low-power mesh networks. By synchronizing the transmitter to the receiver, both motes spend most of

their time with their radio *off*, only turning it *on* when a communication is about to take place. This yields very low radio duty cycles and long network lifetimes. The ideas developed from these protocols are the foundation for standards such as IEEE802.15.4e [2]. The technique presented in this article builds on top of TSMP, TPSN and IEEE802.15.4e. Instead of always re-synchronizing at a "worst case" rate, Adaptive Synchronization lowers the radio's duty cycle even more by modeling the clock drift between neighbor motes, and reducing the effective drift rate by making small corrections that do not require the radio to be powered. This results in a lower idle duty cycle, thus a longer network lifetime. Recently, Kerkez [14] investigated how to achieve tight clock synchronization without a precise clock source. This allows for a board manufacturer to swap an (expensive) crystal for a (cheap) on-chip oscillator. Similar to our approach, Kerkez proposes to account clock drift to be able to correct clock misalignment by modifying slot phase and duration, but requires motes to communicate at every time frame. Very interesting approaches have been published recently, Medina et.al. [15] present a synchronization algorithm based on an estimation of the relative clock drift of each node, their scheme uses periodic beacons to determine clock drift with respect to a global clock shared by all the motes in the network. Besides, in [16] the same authors present a theoretical error characterization of clock drift in similar scenarios. Adaptive synchronization presented in this article takes a similar approach and complements estimated clock drift with periodic temperature monitoring in order to compensate temperature effects on the drift rate of crystals. Liu et.al. [17] uses a Kalman filter to compensate clock drift, their scheme named AdaSynch shows how clock drift can be modeled and fit to a certain value using different Kalman filters. The important aspect is the deep analysis of clock drift that authors carry on a set of 100 telosb motes which can be further used to improve the Adaptive synchronization estimation scheme. Temperature compensation schemes have been studied recently by Brunelli et.al. [18]. Authors propose to divide operation in two phases, a one time calibration where nodes learn the drifting pattern and a operation phases where nodes apply the learnt pattern according to the temperature reading. Adaptive Synchronization composes a temperature synchronization scheme with an adaptive synchronization which benefits from the fact that no calibration is needed and that temperature readings can be spaced in time so energy consumption is not compromised.

III. SYNCHRONIZING IN IEEE802.15.4E

In an IEEE802.15.4e Time Synchronized Channel Hopping (TSCH) network, communication happens in time slots. A network-wide schedule instructs each mote what to do in each slot. The schedule is built to satisfy the throughput and delay requirements of the different traffic flows. A time slot is long enough for the transmitter to transmit the longest possible packet, and for the receiver to send back an acknowledgment indicating reception.
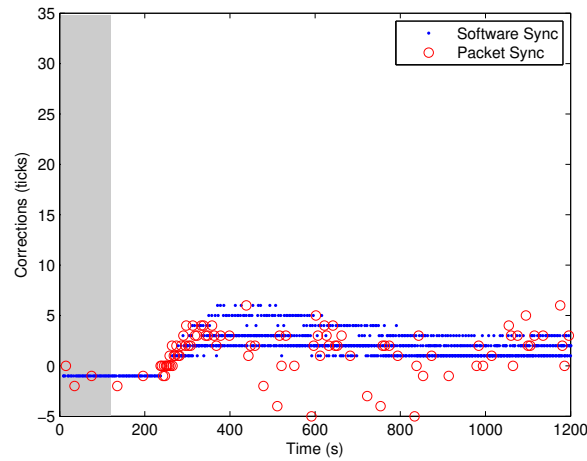
All motes in an IEEE802.15.4e network are synchronized, i.e. neighbor motes are desynchronized by at most a *guard*
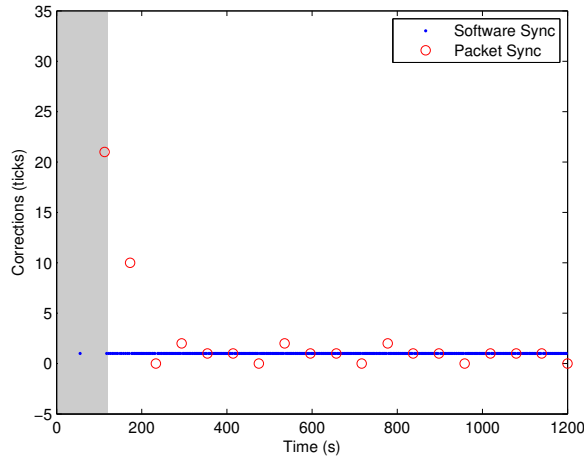
(a) Indoor setup: the temperature is constant.

(b) Outdoor setup: the temperature varies by $15°C$ over the course of the experiment.

(c) Oven setup: the temperature varies by $55°C$ over the course of the experiment.

Fig. 2. Corrections are applied to the clock of a mote which uses Adaptive Synchronization to track the drift with a neighbor mote. Packet-based synchronization, which happens every 60s or when mote temperature has changed $2°C$ since last packet sync, whichever comes first, causes the offset (red circles) to be applied; software corrections are used to further track the calculated drift in-between packet-based synchronizations. The same experiment is conducted in three environments with different temperature variations. Corrections are expressed in $32kHz$ clock ticks.

time, typically $1ms$. The network schedule indicates the "time parents" of each mote, the neighbor motes to which this mote needs to keep synchronized. This results in a synchronization directed acyclic graph. Once a mote has been assigned its time parents, it needs to ensure its clock never drifts by more than a guard time with respect to its time parents.

IEEE802.15.4e is designed to cope with clock drift. It includes timing information in *all* packets, to allow for two neighbor motes to re-synchronize to one another whenever they communicate. When a transmitter transmits a packet, the receiver timestamps the instant at which it receives the packet, and indicates the offset between that measured time and the theoretical reception time `TsTxOffset` in its acknowledgment. As a result, upon each communication, both the transmitter and the receiver know how de-synchronized they are. Depending on which mote is the time parent, either the transmitter or the receiver aligns its clock to the other.


Fig. 3. The GINA mote.

After re-synchronizing, the clocks of both motes are perfectly aligned.

When data packets flow through the network, keeping synchronized comes at no extra cost[3]. When the network sits idle (no data packets are flowing), motes send empty data packets periodically to keep synchronized. The period at which these "keep-alive" packets are exchanged is the object of this article. We call "idle duty cycle" the portion of the time a mote's radio is on just to keep synchronized, when the network sits idle.

Let's take a typical case where a mote has 2 time parents, and 2 children. We assume the mote runs IEEE802.15.4e with a $1ms$ guard time, and is equipped with a $30ppm$ crystal. As per the calculation above, it can expect, each $16s$, to participate in 4 re-synchronizations: to both parents and both children. Each re-synchronization involves exchanging the keep-alive packet (around 15 bytes) and the acknowledgment (around 15 bytes). At $250kpbs$, and taking into account radio startup and turnaround times, this translates to a radio on time of around $2ms$ per synchronization. To participate in all synchronizations, it will have its radio on for around $8ms$ each $16s$, or an idle radio duty cycle of 0.050%.

The goal of the Adaptive Synchronization technique presented in Section IV is to bring this idle duty cycle down by extending the period between two re-synchronization between neighbor motes.

IV. ADAPTIVE SYNCHRONIZATION

Adaptive Synchronization allows a mote to track the drift rate to its neighbor, which it then uses to increase the time between keep-alive messages.

This is done by measuring the *effective* clock drift between two consecutive keep-alive messages, and then using that rate to make periodic "software" adjustments (which do not require communication) in-between "packet-based" re-synchronizations.

At each packet synchronization, the timestamping of the packets indicates to a mote how offset its clock is with respect to its neighbor's. It can calculate the experimental clock drift rate r_{exp} by dividing this offset ε by the duration since the previous packet-based synchronization ΔT, as shown in (2).

$$r_{exp} = \frac{\varepsilon}{\Delta T} \qquad (2)$$

If a mote uses (2) and determines it is fast with respect to a neighbor, it will periodically adjust its own clock in order to "slow down". It can do so by periodically adding a clock tick. This allows it to track its neighbor's clock, thereby reducing the drift. Since the correction is not perfect and the drift rate changes (e.g. with temperature), packet-based re-synchronizations is still needed, but can be less frequent.

Fig. 2(a) shows experimental results gathered on a pair of GINA motes [19] running in a temperature controlled environment. The blue dots show the periodic software corrections applied to the mote's clock; the red circles represent the

[3]Strictly speaking, the time offset between transmitter and receiver is encoded as a 2 byte signed value, which is added to the acknowledgment. Transmitting this extra field costs some energy.


Fig. 4. Correction interval is the interval between each type of sync. For Software Syncs, which in this case has a consistent correction value of a single clock cycle, the interval is closely related to the clock drift rate. This figure shows an overlay of correction intervals for two mote-pairs of different drift rates.

offset indicated by each packet-based synchronization (which is pre-set to happen every $60s$). After a learning period – during which the mote measures the drift rate – the software corrections track the drift, and the offset measured by the packet-based synchronization drops from 22 ticks ($671\mu s$) to 3 ticks ($91\mu s$). This corresponds to an effective drift dropping from $11ppm$ to $1.5ppm$.

Temperature causes the clock drift rate to change. To account for temperature changes, Adaptive Synchronization records the temperature during the most recent packet-sync event and then measures the temperature periodically. If the difference between the recorded temperature and any measurement thereafter is larger than φ, a keep-alive message is triggered to re-synchronize and determine the new effect drift rate. The threshold φ – set to $2°C$ in our experiments – can be tuned to match the crystal's temperature sensitivity and the application requirements.

The other sources of clock drift (supply voltage and crystal aging) change much more slowly and are unlikely to be noticed even with extended keep-alive packet intervals. For this reason we did not make any attempt to measure or correct for these sources.

There is an error associated with the experimental clock drift rate that increases as the packet interval decreases: short intervals make for inaccurate predictions. This could be an issue if a short message interval were immediately followed by a long one. To prevent such a situation, the interval between keep-alive packets is controlled so that a short interval is followed by one twice as long, then four times as long, etcetera, until the target interval length is reached. In our experimental setting we started at a period of $5s$, increasing it up to $60s$. This is similar to TCP's "slow start" [20].

V. EVALUATION

In this section, we experimentally explore the trade-off between synchronization accuracy and energy consumption,

for different crystal oscillator configurations and temperature conditions.

All experiments are conducted with GINA motes [19] (shown in Fig. 3), which feature a Texas Instruments MSP430 micro-controller and an Atmel AT86RF231 IEEE802.15.4 radio chip.

Data was collected at each synchronization event, whether packet or software based. Clock correction amount (in units of ticks), time of correction, a unique sync event ID, temperature, and whether the correction was packet or software based were all recorded.

A. Impact of Manufacturing Differences

To isolate crystal manufacturing, all experiments in this section are conducted with motes operating in the same temperature-controlled environment, corresponding to the indoors case shown in Fig. 2(a).

By letting the motes synchronize without Adaptive Synchronization, we can measure the drift rate. We cherry-pick two pairs of motes: motes H_1 and H_2 exhibit a *high* clock drift of $11ppm$; motes L_1 and L_2 have a *low* effective clock drift of $4ppm$. The motes in both pairs run the same firmware, and are configured to exchange a keep-alive every $60s$.

Fig. 4 depicts the interval between two synchronizations. Because Adaptive Synchronization speeds up synchronization when the drift is large, the interval between software synchronization events is smaller for the (H_1,H_2) $11ppm$ pair than for the (L_1,L_2) $4ppm$ pair. The frequency of synchronization is proportional to the clock drift, so in this case the $11ppm$ pair makes corrections about 2.75 time as frequently as the $4ppm$ pair

B. Impact of Temperature

A pair of motes is setup to send keep-alive messages to keep synchronized using Adaptive Synchronization; no application data is exchanged. To measure the impact of temperature variation the same experiment is repeated in the following cases:

- In the **indoor** case, the pair of motes is placed on a desktop in an office environment. The experiment is conducted over the course of 20 hours, during which the temperature is constant.
- In the **outdoor** case, one mote is placed indoors, the other one outdoors. The experiment is conducted over the course of 12 hours, during which the temperature drops by $15°C$.
- In the **oven** case, one mote is placed at a constant $20°C$, the other one is placed in an oven. From $20°C$, the temperature of the oven is brought to $75°C$ over the course of 2 minutes, then down to $55°C$ over 10 minutes.

In each case, Adaptive Synchronization is running with a preset keep-alive interval of $60s$. As detailed in Section IV, sudden changes in temperature trigger extra keep-alive messages.

Fig. 5 shows the total accumulated correction conducted by Adaptive Synchronization in the three cases. The total amount


Fig. 6. The interval between time correction events in the oven setup scenario. Packet-based corrections are separated from software-based corrections. At around 250s, the temperature is increased using the oven. The slow start process can be seen when temperature reaches a steady state at around 350s to 400s.

that Adaptive Synchronization corrected is equivalent to how the two motes would drift when corrections are not applied, therefore, Fig. 5 presents as well the accumulated drift between two motes in the presented scenarios. A change in temperature alters the accumulation rate. In the indoor case, the value of accumulated correction over time is increasing linearly due to stable temperature. There is slow variation in the accumulation rate of correction in the outdoor case, caused by the slow temperature drift during the transition between day and night. In the oven case, the green line shows the sharp temperature change when the oven is switched on.

Figure 6 shows the interval between time correction events in the oven setup. As the temperature increases, it can be seen the packet synchronizations happen frequently at around $250s$ as the temperature rapidly increases. When temperature stabilizes, Adaptive Synchronization starts the slow start process to achieve the maximum keep alive interval.

C. Energy Consumption Discussion

Fig. 7 shows a direct comparison between synchronizing with and without Adaptive Synchronization. Because drift is not compensated when not using Adaptive Synchronization the offset is large at each packet-based synchronization. In Fig. 7, this is around 20 ticks, or a drift of $11ppm$. When using Adaptive Synchronization, drift is compensated, and the offset triggered by each packet-synchronization drops to 2 ticks, or $1ppm$.

This means that the guard time can be much smaller. In this case, without Adaptive Synchronization, the guard time must be at least $660\mu s$ long (i.e. the clock drift in 60s); it can be $60\mu s$ when using Adaptive Synchronization, using the same $60s$ keep-alive interval. This translates into a 90% reduction of the energy spent by the mote idle listening for packets.
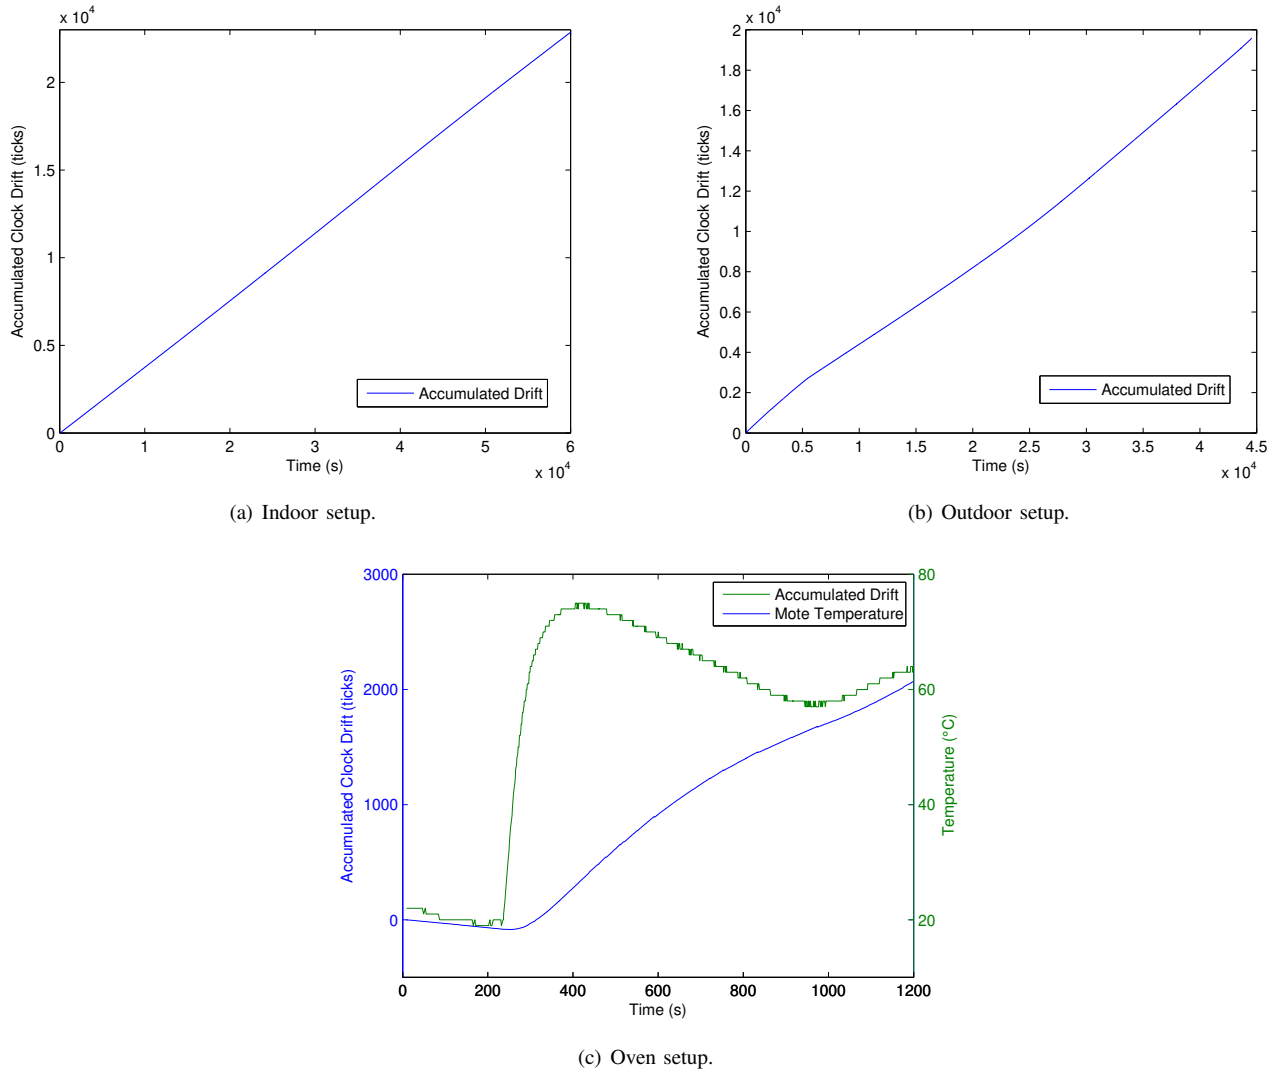
(a) Indoor setup.

(b) Outdoor setup.

(c) Oven setup.

Fig. 5. Accumulated Clock Drift for a mote pair. This is the accumulated value of corrections made by both software adjustments and packet-based re-synchronizations which shows how two motes would drift when corrections are not applied. The figures show how the drift rate, ie. the slope, changes with temperature.

VI. CONCLUSION

Time Synchronized Channel Hopping networks such as those defined by the IEEE802.15.4e standard require tight synchronization of their motes, which limits the lowest achievable duty cycle and hence the minimum energy consumption of the network. Clock drift is attributed to crystal manufactoring differences and temperature, mainly. TSCH uses periodic keep-alive messages to re-synchronize neighbor motes in the absence of application traffic.

This article presents a simple Adaptive Synchronization technique which can either reduce the keep-alive interval, or shorten the guard time. In both cases, this translates directly in a reduction of the idle duty cycle. Experimental results show an idle duty cycle reduced by a factor of 10, with the ability to maintain synchronization even in rapidly varying temperature settings. The technique is implemented in the IEEE802.15.4e MAC layer of the OpenWSN protocol stack, and evaluated on off-the-shelf motes.

Fig. 7. Correction values for the $11ppm$ mote-pair running Adaptive Synchronization versus packet-only correction.

REFERENCES

[1] *WirelessHART Specification 75: TDMA Data-Link Layer*, HART Communication Foundation Std., Rev. 1.1, 2008, hCF_SPEC-75.

[2] *802.15.4e-2012: IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer*, IEEE Std., 16 April 2012.

[3] *802.15.4-2006: IEEE Standard for Information technology, Local and metropolitan area networks, Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs)*, IEEE Std., 2006.

[4] E. Toscano and L. L. Bello, "Multichannel Superframe Scheduling for IEEE802.15.4 Industrial Wireless Sensor Networks," *IEEE Trans. Industrial Informatics*, vol. 8, no. 2, pp. 337–350, 2012.

[5] T. Watteyne, A. Mehta, and K. S. J. Pister, "Reliability Through Frequency Diversity: Why Channel Hopping Makes Sense," in *Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, October 2009.

[6] T. Watteyne, X. Vilajosana, B. Kerkez, F. Chraim, K. Weekly, Q. Wang, S. D. Glaser, and K. Pister, "OpenWSN: a Standards-based Low-power Wireless Development Environment," *Transactions on Emerging Telecommunications Technologies*, vol. 23, no. 5, pp. 480–493, 2012.

[7] W. C. Lindsey, F. Ghazvinian, W. Hagmann, and K. Dessouky, "Network Synchronization," *Proceeding of the IEEE*, vol. 73, no. 1, pp. 1445–1467, Oct 1985.

[8] J.-H. Chen and W. Lindsey, "Mutual Clock Synchronization in Global Digital Communication Networks," *European Transactions on Telecommunications*, vol. 7, no. 1, pp. 25 –37, January-February 1996.

[9] Z. Zhang, H. Kayama, and C. Tellambura, "New Joint Frame Synchronisation and Carrier Frequency Offset Estimation Method for OFDM Systems," *European Transactions on Telecommunications*, vol. 20, no. 4, pp. 413–430, 2009.

[10] C. H. Rentel and T. Kunz, "Network Synchronization in Wireless Ad Hoc Networks," Carleton University, Tech. Rep., July 2004.

[11] M. M. H. P. van den Heuvel, R. J. Bril, and J. J. Lukkien, "Transparent synchronization protocols for compositional real-time systems." *IEEE Trans. Industrial Informatics*, vol. 8, no. 2, pp. 322–336, 2012.

[12] K. S. J. Pister and L. Doherty, "TSMP: Time Synchronized Mesh Protocol," in *International Symposium on Distributed Sensor Networks (DSN)*. Orlando, Florida, USA: IASTED, November 2008.

[13] S. Ganeriwal, R. Kumar, and M. B. Srivastava, "Timing-sync Protocol for Sensor Networks," in *International Conference on Embedded Networked Sensor Systems (SenSys)*. New York, NY, USA: ACM, 2003, pp. 138–149.

[14] B. Kerkez, "Adaptive Time Synchronization and Frequency Channel Hopping for Wireless Sensor Networks," Master's thesis, University of California, Berkeley, EECS Department, June 2012.

[15] C. Medina, J. C. Segura, and A. De La Torre, "Accurate time synchronization of ultrasonic tof measurements in ieee 802.15.4 based wireless sensor networks," *Ad Hoc Netw.*, vol. 11, no. 1, pp. 442–452, Jan. 2013. [Online]. Available: http://dx.doi.org/10.1016/j.adhoc.2012.07.005

[16] C. Medina, J. C. Segura, and A. de la Torre, "A synchronous tdma ultrasonic tof measurement system for low-power wireless sensor networks," *Instrumentation and Measurement, IEEE Transactions on*, vol. PP, no. 99, pp. 1 –13, 2012.

[17] Q. Liu, X. Liu, J. L. Zhou, G. Zhou, G. Jin, Q. Sun, and M. Xi, "Adasynch: A general adaptive clock synchronization scheme based on kalman filter for wsns," *Wirel. Pers. Commun.*, vol. 63, no. 1, pp. 217–239, Mar. 2012.

[18] D. Brunelli, D. Balsamo, G. Paci, and L. Benini, "Temperature compensated time synchronisation in wireless sensor networks," *Electronics Letters*, vol. 48, no. 16, pp. 1026 –1028, 2 2012.

[19] A. Mehta and K. Pister, "WARPWING: A Complete Open-Source Control Platform for Miniature Robots," in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, 2010.

[20] M. Allman, V. Paxson, and E. Blanton, "TCP Congestion Control," RFC 5681 (Draft Standard), Internet Engineering Task Force, September.