ICN over TSCH: Potentials for Link-Layer Adaptation in the IoT

Oliver Hahm
INRIA
oliver.hahm@inria.fr

Cédric Adjih
INRIA
cedric.adjih@inria.fr

Emmanuel Baccelli
INRIA
emmanuel.baccelli@inria.fr

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

ABSTRACT

In this abstract, we start discussion on (a) leveraging ICN communication patterns to dynamically optimize the use of TSCH (Time Slotted Channel Hopping), a wireless link layer technology increasingly popular in the IoT. Through a series of experiments on FIT IoT-LAB interconnecting typical IoT hardware, we find that our proposal is fully robust against wireless interference, and almost halves the energy consumed for transmission when compared to CSMA. Most importantly, our adaptive scheduling prevents the time-slotted MAC layer from sacrificing throughput and delay.

CCS Concepts

• **Networks** → **Network experimentation; Naming and addressing; Cyber-physical networks;**

Keywords

IoT, NDN, TSCH, 802.15.4e, name-based routing, adaptive forwarding

1. INTRODUCTION

The typically unreliable and fluctuating nature of wireless communication in the IoT has a strong impact on the functionality of an ICN layer. This motivates the search for an adaptive link layer technology, which allows appropriate cooperative use of the radio. A promising candidate is Time Slotted Channel Hopping (TSCH) [2], which replaces CSMA with a reservation-based MAC protocol, combining TDMA with frequency hopping.

TSCH can drastically increase the reliability of transmission [3] thereby guaranteeing a fixed throughput and maximum latency even at high traffic load—if a proper schedule exists. However, an a priori schedule requires thorough understanding of future traffic flows, which is infeasible for most application domains. Furthermore, traffic patterns may vary over time, leading to fluctuating demands that contradict a static schedule.

High-level idea. In this poster, we argue that the schedule for media access can be derived from inherent NDN traffic patterns, namely the request response property in NDN: Data only flows *after* an Interest message was observed. Consequently, static time slots are necessary only to observe Interests. These slots can be short. Then, longer time slots can be reserved dynamically following Interest observations.

We believe that the long-term success of NDN is tightly coupled to identifying clear advantages over IP. The request response pattern is one of the unique properties in NDN. The application of this pattern should be explored further. We start the discussion on the adaption of the TSCH link layer based on NDN traffic pattern by introducing the general idea and concrete schedules (§ 2) and present first promising evaluation (§ 3).

2. TIME SLOTTED CHANNEL HOPPING FOR ICN IN THE IOT

Content distribution in NDN follows a request response pattern with footprint on each hop. Each Pending Interest (PI) matches at most one data chunk of limited size. Hence, in a fully deterministic, lossless setting, each request is answered by a train of up to k data packets within a time frame bound by the (temporal) diameter of the network. For scheduling the wireless, we can interpret an Interest as a predictor of data expected on the reverse path, and conversely can exclude any data arrival in the absence of PI state. We can further exploit the predefined chunk size for fixing the ratio of data per Interest packet in our schedule. Ideally, the arrival of an Interest would trigger the allocation of k slot frames towards the appropriate neighbor at the expected time.

The general idea is a schedule that is partly static and pre-reserved, and partly dynamic and adaptive to the current traffic pattern. For this, we divide the slotframe into three parts, henceforth called subslotframes (SSFs). The first, named SSF_I, is dedicated to statically scheduled Interest propagation. Second, SSF_C is for sending back content chunks on a semi-dynamic schedule. The third, SSF_{Dyn} is fully dynamic, activated to serve increased traffic on dedicated links.

Examples of MAC Configurations. The static schedule and potential forwarding paths are computed beforehand. The static schedule for SSF_I and SSF_C ensures basic connectivity and reserves one cell per link and direction. We use a length of 15 ms for the slot length and 101 slots per slotframe. The remaining cells are left initially unscheduled and

ACM-ICN '16, September 26–28, 2016, Kyoto, Japan

ACM ISBN 978-1-4503-4467-8/16/09…$15.00

DOI: http://dx.doi.org/10.1145/2984356.2985226

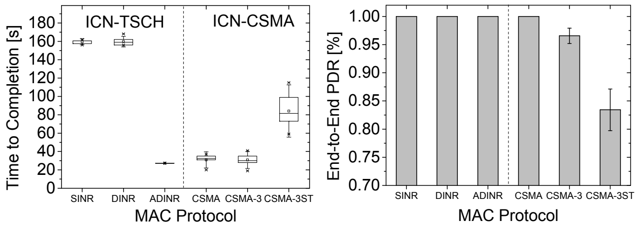

(a) Time to Completion (b) Packet Delivery Ratio

Figure 1: Time to completion and PDR in different configurations for TSCH and CSMA.

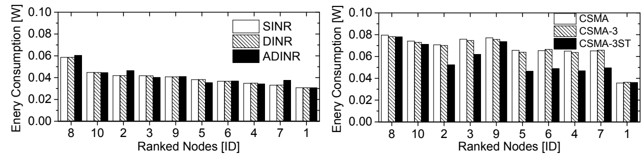

(a) TSCH (b) CSMA

Figure 2: Energy consumption for the different configurations of TSCH and CSMA.

can be reserved in SSF_{Dyn}. The schedule ensures that packets can travel between nodes within one slotframe in SSF_I and SSF_C respectively. Time synchronization between the nodes is based on periodic broadcasting of enhanced beacons in shared cells. We now present four example schemes.

Static Information-centric Networking Reservation (SINR). All scheduled cells in SSF_I and SSF_C are active and the consumer sends out Interests with a constant rate of one Interest per slotframe.

Dynamic Information-centric Networking Reservation (DINR). The scheduled cells in SSF_C are kept initially inactive. Again, the consumer sends out Interests with a constant rate of one Interest per slotframe. As soon as a node A on the path from consumer to producer receives an Interest, it activates its RX cell(s) in SSF_C on the link to the next hop B on the path. Once A receives the corresponding content chunk from B, it deactivates the cell again.

Adaptive Dynamic Information-centric Networking Reservation (ADINR). ADINR is similar to *DINR*, but makes also use of the dynamic part of the schedule SSF_{Dyn}. If a node A receives a certain rate of Interests from one of its neighbors, it activates cells in SSF_{Dyn} in both directions to increase the bandwidth on this link. Conversely, if cells for this link are less frequently used, SSF_{Dyn} cells are deactivated again. In this configuration we increase the rate in which the consumer generates to 15 Interests per slotframe.

3. EVALUATION

Experiment Setup. We conducted our experiments in a ten node multi-hop network, including one consumer and one producer. We compare our proposal with an implementation that runs ICN directly on the link layer, using CSMA as a MAC protocol [1]. Implementation is based on the standard implementation of IEEE 802.15.4e, OpenWSN, and RIOT. The requested content consists of 100 chunks. We generated side traffic in order to create a realistic scenario.

Results. We consider four different metrics: (i) time to completion, (ii) jitter, (iii) end-to-end (e2e) packet delivery ratio (PDR), and (iv) energy consumption.

Since only one Interest and one content chunk can be transmitted per slotframe with *SINR* and *DINR*, the minimum time to completion for fetching 100 content chunks is $\Delta = 100 * T_{SF}$ with T_{SF} being the duration of the slotframe. Our results are only slightly above this minimum (see Fig. 1(a)). Initially, *ADINR* generates more Interests per slotframe than it can send out, but gradually, nodes along the path activate more cells in SSF_{Dyn}. SSF_{Dyn} contains 70 cells, which implies that up to 8 additional links per hop can be scheduled. This leads to a significant improvement of the time to completion in comparison to *SINR* and *DINR*. Concerning jitter, we observe very small standard deviation for SINR, DINR, and ADINR (as expected for a reservation-based MAC) while in comparison, time to completion with CSMA is much less predictable. As shown in Fig. 1(a), CSMA performance depends more heavily on the number of collisions and retransmissions (per link and end-to-end), standard deviation and average increasing substantially if side traffic increases.

As expected with a collision-free TSCH schedule, we observed 100% e2e PDR and almost no link layer retransmissions with *SINR*, *DINR* and *ADINR* (< 5 retransmissions overall). With *CSMA*, Interests are retransmitted as often as required, and thus e2e PDR reaches 100% too, but at the cost of retransmissions and duplicates. On average, we counted more than 130 e2e retransmissions and 25 duplicate chunks that arrived at the consumer. Limiting the number of retransmissions to three (in *CSMA-3*, see 1(b)) decreases the PDR to about 97%, with similar numbers for retransmissions and duplicates. Additional side traffic (in *CSMA-3ST*) decreases the PDR further, with more retransmissions and duplicates.

For energy measurements, we computed the average over all samples in all experiment runs, per node. With TSCH, transceivers switch to sleep mode for all unscheduled slots. Thus, we can see that all nodes consume consistently less power with *SINR*, *DINR*, and *ADINR* than with *CSMA* (see Figure 2). Furthermore, the increased energy consumption in *ADINR* due to a higher duty cycle is leveled out by the fact that the nodes can more quickly return to sleep mode again.

Acknowledgment. This work was partially supported by the German Federal Ministry of Education and Research within the project I3.

4. REFERENCES

[1] E. Baccelli, C. Mehlis, O. Hahm, T. C. Schmidt, and M. Wählisch. Information Centric Networking in the IoT: Experiments with NDN in the Wild. In *Proc. of 1st ACM Conf. on ICN*, pages 77–86, New York, 2014. ACM.

[2] IEEE. Low-Rate Wireless Personal Area Networks, Amendment 1: MAC sublayer. Technical Report IEEE Std 802.15.4e-2012, 2012.

[3] T. Watteyne, V. Handziski, X. Vilajosana, S. Duquennoy, O. Hahm, E. Baccelli, and A. Wolisz. Industrial wireless ip-based cyber–physical systems. *Proceedings of the IEEE*, 104(5):1025–1038, 2016.