

ICNLoWPAN – Named-Data Networking for Low Power IoT Networks

Cenk Gündoğan
HAW Hamburg

cenk.guendogan@haw-hamburg.de

Peter Kietzmann
HAW Hamburg

peter.kietzmann@haw-hamburg.de

Thomas C. Schmidt
HAW Hamburg

t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin

m.waehlich@fu-berlin.de

Abstract—Information Centric Networking is considered a promising communication technology for the constrained IoT, but NDN was designed only for standard network infrastructure.

In this paper, we design and evaluate an NDN convergence layer for low power lossy links that (1) augments the NDN stateful forwarding with a highly efficient name eliding, (2) devises stateless compression schemes for standard NDN use cases, (3) adapts NDN packets to the small MTU size of IEEE 802.15.4, and (4) generates compatibility with 6LoWPAN so that IPv6 and NDN can coexist on the same LoWPAN links. Our findings indicate that stateful compression can reduce the size of NDN data packets by more than 70 % in realistic examples. Our experiments show that for common use cases ICNLoWPAN saves 33 % of transmission resources over NDN, and about 20 % over 6LoWPAN.

Index Terms—Internet of things, low power lossy networks, 802.15.4, header compression

I. INTRODUCTION

The Internet of Things (IoT) has been identified as a potential deployment area for Information Centric Networks (ICN) [1], since infrastructureless access to content, resilient forwarding, and in-network data replication demonstrated notable advantages over the traditional host-to-host approach on the Internet [2]–[5]. Named Data Networking (NDN) [6] has matured to a prominent, widely implemented flavor of ICN that strictly couples data delivery to consumer requests. Recent studies [7] have shown that NDN can outperform CoAP [8] and MQTT-SN [9], the corresponding IP-based data services for the constrained IoT. However, NDN lacks mechanisms for packet adaptation and compression to comply with limitations of wireless links prevalent at the low power IoT edge [10].

Common edge networks in the constrained IoT are built from low power and lossy radios (see "LLN" in [11]) such as IEEE 802.15.4 [12], Bluetooth Low Energy (BLE) [13], or LoRA [14]. Characteristics of LLNs include an unreliable environment, low bandwidth transmissions, and increased latencies. IEEE 802.15.4 admits a maximum physical layer packet size of 127 octets. With 6LoWPAN [15], [16], the IP-world has created a convergence layer that provides appropriate frame encapsulation formats, packet header compression and link fragmentation for IPv6 packets in IEEE 802.15.4 networks. The ICN world has not yet developed corresponding features for constrained environments.

ISBN 978-3-903176-16-4 © 2019 IFIP

With this paper, we close the gap of LoWPAN convergence for NDN. We leverage its potential of stateful forwarding for eliding names on paths and design highly efficient compression primitives that outperform 6LoWPAN. Our evaluations reveal significant gains in packet reduction, energy consumption, and reliability. In addition to stateful and stateless compression, we also contribute a fragmentation scheme as well as a framing compatible to the 6LoWPAN techniques. Real implementations under RIOT OS [17], [18] and experimentation on a testbed of current IoT hardware demonstrate the feasibility, robustness and energy efficiency of our approach.

The remainder of this paper is organized as follows. The subsequent Section II discusses the problem of NDN LoWPAN adaptation and related work. In Section III, we introduce our ICNLoWPAN convergence layer and detail out IPv6-ICN coexistence, on-link fragmentation, and compression. A thorough evaluation of the compression benefits follows in Section IV. Finally, we conclude with an outlook in Section V.

II. PROBLEM SPACE AND RELATED WORK

The Internet of Things inherently connects numerous devices of substantial heterogeneity. In this work, we focus on deployment use cases that bundle low-end and battery-operated microprocessors in wireless networks, where packet transmission distinctly dominates power consumption. The challenges we face in such scenarios are manifold and range from limited MTUs, lossy links and mobility to link layers that lack basic protocol features, such as frame encapsulation formats (c.f. EtherTypes in Ethernet).

NDN couples name-based routing from TRIAD [19] with stateful forwarding from DONA [20] and seamlessly leverages in-network caching on the forwarding plane. The fundamental request-response semantic on the network layer of NDN requires an *Interest* message and a returning *data* message. Both message types utilize flexible Type–Length–Value (TLV) header fields to allow for generic and extensible packet formats to the cost of space efficiency. Name TLVs are essential to NDN and thus always appear in *Interest* as well as in *data* messages. Depending on the naming scheme, human-readable Name TLVs make up the largest part of a request, and also of a response for many IoT use cases. We explore related work that copes with strict message length limitations using header

compression and fragmentation in IP networks first, and then discuss proposed solutions for NDN.

IPv6 mandates a minimum MTU of 1280 bytes for each link and thus precludes a proper IPv6 operation in low power networks with small-sized MTUs. The IETF designed and extends a set of protocols for constrained IoT deployments where the 6LoWPAN convergence layer is an integral part of. It is situated below the network layer and provides packet encapsulation, stateless and stateful header compression as well as a protocol independent link fragmentation scheme. A generic header compression (GHC) [21] extends 6LoWPAN with an LZ77 flavored approach to deal with headers and header-like payloads that are not covered by the 6LoWPAN compression specification. While 6LoWPAN proves necessary for an interoperable and interconnected host-centric IoT, the same challenges remain open for information-centric IoT deployments.

Shang *et al.* [22] proposed a lightweight link fragmentation scheme that prepends a 3-byte fragmentation header to each NDN fragment to allow for messages larger than the limiting MTU of IEEE 802.15.4. This custom header further supports a minimal protocol identification by distinguishing normal NDN messages from fragmented messages. However, this protocol encapsulation collides with the 6LoWPAN encapsulation and thus disregards interoperability with 6LoWPAN, especially in multi-interface and multi-protocol deployments.

Another approach was presented by Mosko *et al.* [23], which extends NDN with a new message type that encapsulates each message fragment. It also adds complexity to the state machine for each peer to initialize fragment sequence numbers and perform corrective actions in case of drifting sequences after packet loss. A similar approach is the NDN Link Protocol (NDNLP) [24], which features fragmentation and reassembly as well as ARQ mechanisms. It provides packet encapsulation to distinguish between normal NDN messages and link acknowledgments. The last two approaches add overhead in terms of memory consumption and error-control messages which is a disadvantage for low power use cases.

A secure fragmentation for content-centric networks that does not rely on hop-by-hop reassembly and therefore decreases latency was derived by Ghali *et al.* [25]. Each fragment is securely signed according to NDN semantics and can be cached on intermediate routers. The authors propose a new ContentFragment message type that includes a Name TLV for forwarding purposes. Since NDN names are theoretically of unlimited length, duplicating names for each ContentFragment message adds a significant overhead, which naturally is controversial in constrained IoT networks. Furthermore, sporadic disruptions and mobility in LLNs do not guarantee a successful handover of each individual fragment to complete the reassembly.

Yang *et al.* [26] focused on bandwidth reduction and an improved storage utilization by translating long names into short names for local communication. In this regard, a sensor node registers a prefix at a sink node and receives a shorter

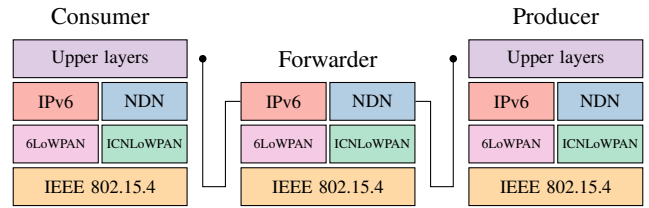


Fig. 1: Stack traversal in a 6LoWPAN and ICNLoWPAN.

prefix, e.g., a hash as a name replacement. Ingress messages that traverse the sink node are updated to include the short name and egress messages respectively are updated to include the long name. The proposed local name translation is transparent to nodes outside the local IoT network. However, we argue that a centralized approach to handle name translations with registration and cancellation procedures adds complexity that limits the scalability in large deployments.

In the following section, we concentrate on a fully distributed convergence layer that preserves compatibility with IP LoWPAN, but takes advantage of the information-centric characteristics to obtain a generically applicable, efficient though uncomplex compressive encoding named ICNLoWPAN.

III. ICNLoWPAN

ICNLoWPAN provides a convergence layer that maps ICN packets onto constrained link layer technologies to enable pure NDN deployments without running as an overlay on top of IP. Our convergence layer includes features such as link fragmentation, protocol separation on the link layer level as well as stateless and stateful header compression mechanisms. Fig. 1 shows the overall network stack of a 6LoWPAN and an ICNLoWPAN deployment in parallel for a consumer, a forwarder, and a producer. Both convergence layers are situated between the link layer and the actual network layer, such that each message traverses them. This allows for a transparent operation without the need for modifications on the network layer.

A. 6LoWPAN Dispatching Framework

6LoWPAN defines a dispatching framework in [15], where each frame is prepended with a 1-byte dispatch type and a possible dispatch header. Several dispatch types exist already, e.g., for stateless IPv6 compression, a mesh header for mesh-under routing purposes, link fragmentation, and extensions to expand the limited universe of possible dispatch types. One of those extensions is the 1-byte page switch dispatch [27], which arranges dispatch types into 16 pages and signals a context switch to the packet parser to choose the proper page before interpreting subsequent dispatch types.

ICNLoWPAN integrates into the 6LoWPAN dispatching framework by defining four new dispatch types for Interest and data messages that are either compressed or uncompressed. Since page 0 and page 1 are already reserved for 6LoWPAN usage, we allocate these dispatches from page 2 to allow for coexistence with 6LoWPAN deployments. A prepended page

switch dispatch before the very first ICNLoWPAN dispatch is thus necessary as an indicator to the dispatch parser.

A typical ICNLoWPAN message encapsulated in an IEEE 802.15.4 frame is shown in Fig. 2. RFC4944 dispatches are optional and may include all dispatch types defined in [15]. Note the 1-byte page 2 dispatch before the first ICNLoWPAN dispatch. To switch back to 6LoWPAN dispatches after an ICNLoWPAN dispatch, another page switch dispatch to page 0 or 1 is necessary.

A major benefit of reutilizing the 6LoWPAN dispatching framework is to share a common code base for dispatch handling. Notably for multi-interface and multi-protocol deployments that make use of IPv6 and NDN simultaneously, using the same code components in resource-constrained devices is exceptionally valuable for minimizing RAM and ROM requirements.

B. Fragmentation

Reusing the 6LoWPAN dispatching framework enables ICNLoWPAN to seamlessly benefit from the protocol independent link fragmentation scheme defined in [15]. It is thus possible to fragment large NDN messages to fit the limited maximum physical packet sizes of low power link layers, such as 127 bytes for IEEE 802.15.4.

Practically, a fragmented NDN message includes a 4-byte fragmentation dispatch header that lists the original datagram size and a datagram tag to identify fragments of differing packets. Subsequent fragments further include an additional 1-byte datagram offset of the payload in the dispatch header. Fragments are reassembled on the next hop and passed to the NDN network stack as typical NDN Interest or data messages. The 6LoWPAN fragmentation scheme does not define ARQ mechanisms to recover lost fragments, but rather relies on corrective actions of the link layer. This allows for implementations with minimal memory footprints.

C. Stateless Compression

ICNLoWPAN defines a stateless header compression scheme with the main purpose of reducing header overhead of NDN packets. This is of particular importance for link layers with small MTUs and for increasing energy conservation of battery-operated and wirelessly connected devices. Corresponding dispatch headers in the ICNLoWPAN packet provide the rule set for decompressing NDN messages before handing the packet over to the NDN network stack.

1) *TLV Compression*: The NDN header format is solely composed of TLV fields to encode header data. The advantage of TLVs is a native support of variable-sized data. The main disadvantage of TLVs is its verbosity that results in two extra bytes for each header field from storing the *type* and *length* of the encoded data.

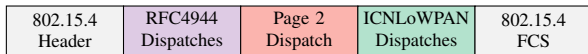


Fig. 2: IEEE 802.15.4 encapsulated ICNLoWPAN message.

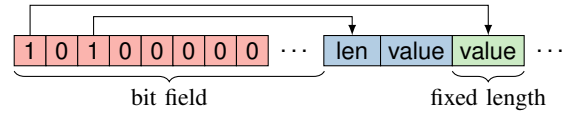


Fig. 3: Eliding *type* and *length* fields using compact bit fields.

The stateless header compression scheme of ICNLoWPAN makes use of compact bit fields to indicate the presence of optional TLVs in the uncompressed packet as illustrated in Fig. 3. Each *type* that is present in the bit field is thus elided from the actual TLV representation, which translates to a reduction from 1 byte to 1 bit. Further compression is achieved with eliding the *length* of TLVs that either represent fixed-length header data, or where the length can be assumed from surrounding TLVs. We also achieve smaller encodings by specifying sane default configurations for IoT use cases.

2) *Name Compression*: A Name TLV is substantial to NDN messages and usually consists of several name components, each of variable size. An Interest message essentially carries the Name TLV and a data message returns either the same TLV, or a more specific variant with an equal prefix. The NDN TLV encoding requires at least two bytes for each name component (*type* + *length*) and an extra two bytes for the outer-most Name TLV. The TLV overhead for a name is displayed in Eq. 1, where $|c|$ is the number of components.

$$TLV\ overhead = 2 + 2 \cdot |c| \quad (1)$$

ICNLoWPAN provides a compression scheme for Name TLVs that drastically reduces the TLV overhead of each nested component. This compression encodes length fields of two consecutive component TLVs into one byte, using 4 bits each as displayed in Fig. 4. This process limits the length of a component TLV to 15 bytes. To further elide the outer-most *length* field of the name, this scheme utilizes a stop marker. For an odd number of components, the stop marker is encoded into the least significant 4 bits of the current *length* byte. On an even number of components, the full *length* byte is already occupied with two name components. In this case, a stop byte is appended to the last component as shown in Fig. 4.

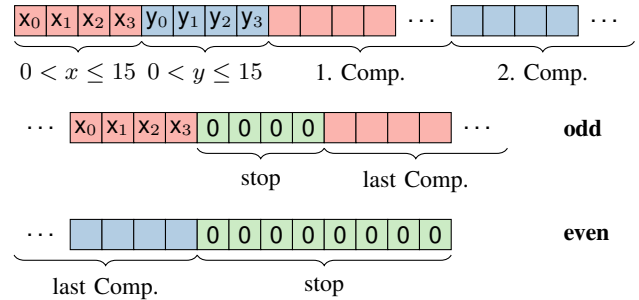


Fig. 4: Stateless name compression and stop marker for odd and even number of name components.

Compressed names yield a significantly lower TLV overhead as displayed in Eq. 2, where $|c|$ again is the number of

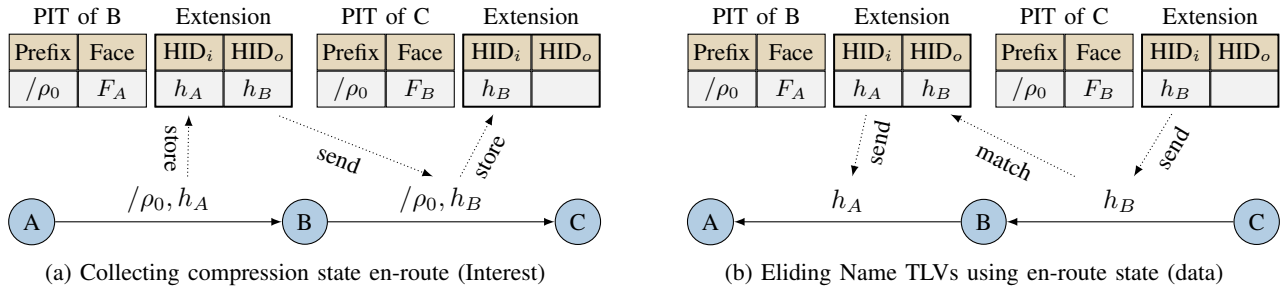


Fig. 5: Stateful header compression using en-route forwarding state.

components. The ceil operator handles both cases of odd and even for $|c|$.

$$TLV \text{ overhead} = \left\lceil \frac{|c| + 1}{2} \right\rceil \quad (2)$$

The total TLV overhead reduction for names thus follows from Eq. 1 and Eq. 2 and is a function of $|c|$: $\lceil 1.5 \cdot |c| \rceil + 1$.

D. Stateful Compression

ICNLoWPAN further employs two stateful compression schemes to enhance size reductions. These mechanisms rely on shared contexts that are either distributed and maintained in the whole LoWPAN, or are generated and maintained on-demand for a particular Interest-data path. Our stateless and stateful compressions are applied in succession to produce huge compression savings, which we show in Sec. IV.

1) *LoWPAN-local State*: A context identifier (CID) is a 1-byte number that refers to a particular conceptual context between networked devices and may be used to replace frequently appearing information, like name prefixes, suffixes, or meta information, such as an Interest lifetime. This allows for a reduction of potentially long data to a single byte. Shared context has to be initially distributed on compile-time or dynamically maintained on run-time in order for a device to properly encode and decode NDN messages.

The convergence layer replaces header fields of outgoing Interest and data packets with CIDs maintained in a CID state table. Context identifiers follow the last ICNLoWPAN dispatch, while the most significant bit of a CID signals the presence of a subsequent CID. On reception, the original packet is restored and passed to the network layer.

2) *En-route State*: An NDN Interest requests data by a name or a prefix which is then returned in the corresponding data packet. This duplication generates large overhead in particular for long names. To deduplicate we make use of ephemeral 1-byte HopIDs that replace the name in data responses and link them to entries of the *Pending Interest Table (PIT)*. The PIT is a fundamental component of NDN that hop-wise matches the returning responses to open requests by name. HopIDs must be unique within the local PIT and only exist during the lifetime of a PIT entry. We extend the PIT by two new columns to manage these HopIDs. HID_i for inbound HopIDs and HID_o for outbound HopIDs as visualized

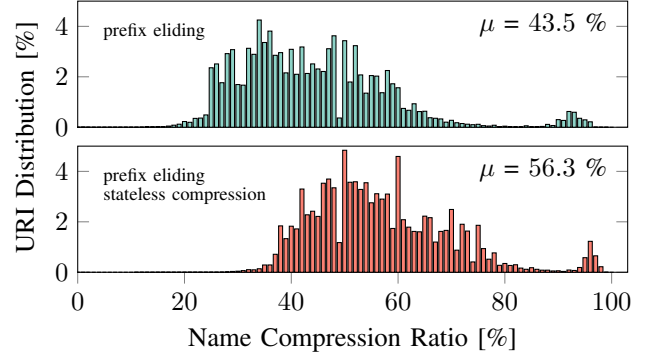


Fig. 6: Distribution of percental name compression ratios.

in Fig. 5. We emphasize that even though PIT entries are extended by two additional bytes, the overall RAM consumption on link commitment reduces due to smaller packets and corresponding message buffers.

Before sending an Interest, a node generates a HopID and stores it in the local HID_o column. This Interest then includes the generated HopID along with the name. On the next hop, the HopID is extracted from the Interest and stored in the HID_i column of the respective PIT entry. The forwarder then generates a new HopID, stores it in the HID_o column of the particular PIT entry, and puts this HopID into the Interest message before it is forwarded to the next hop. This process is repeated for each hop until the request can be satisfied with the corresponding response as displayed in Fig. 5a.

The producer of a returning data message reverses this process by obtaining a HopID from the HID_i column of a PIT entry and encodes it into the response message. If the returning name equals the original name, then it is fully elided. Otherwise, the distinct suffix is included along with the HopID. When a response is forwarded, the contained HopID is extracted and used to match against the correct PIT entry by performing a lookup on the HID_o column. The HopID is then replaced with the corresponding HopID from the HID_i column before forwarding the response, as visualized in Fig. 5b.

E. Name Compression Performance

Our proposed stateless and stateful name compression mechanisms are of simple nature and do not exhaust scarcely available CPU resources. The compression ratio is highly

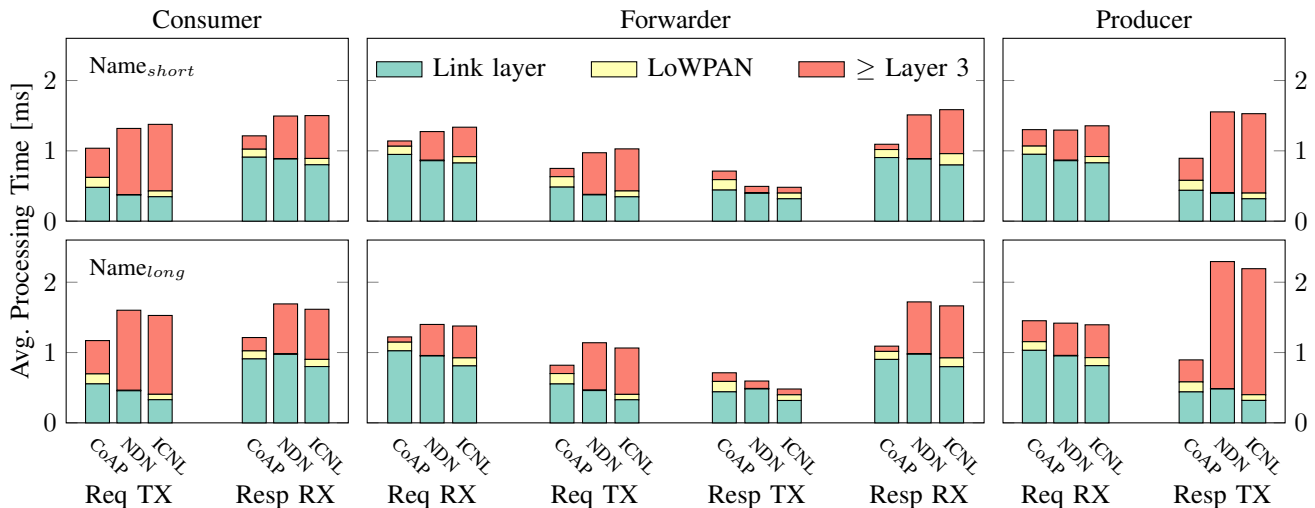


Fig. 7: Intra-stack average processing times for CoAP, NDN, and ICNLoWPAN

dependent on 1) the number of name components as they dictate the overall TLV overhead and 2) the length of the prefix omitted from the name. To quantify the effects, we analyze the average name compression performances in Fig. 6 for $356 \cdot 10^6$ real-world URI paths obtained from the WWW, where each component does not exceed 15 bytes. Before applying our compression scheme, we encode all names as NDN Name TLVs. First, we elide the hostname from each URI as part of our stateful compression in Sec. III-D. This yields an average compression ratio of 43.5 %. Second, we apply the stateful name compression described in Sec. III-C2 to reduce the TLV overhead and observe a compression ratio of 56.3 % on average. We believe that typical names in a low power IoT edge network will yield similar, if not better, compression performances.

IV. EVALUATION

A. Experiment Setup

1) *Protocol Comparison*: We consider two different NDN deployments for low power IoT networks. In the first experiment, we run NDN directly on top of IEEE 802.15.4. Due to the small MTU of this particular link layer, we limit packet sizes to a maximum of 100 bytes. In the second experiment, NDN runs on top of ICNLoWPAN with activated stateless and stateful compression. We further compare our NDN setups with a typical 6LoWPAN operation that uses UDP as a transport protocol and CoAP as an application protocol. We specifically compare against the GET method of CoAP as it provides a request-response pattern analog to NDN. We deploy our devices in two different network configurations.

a) *Single-hop*: A consumer device has managed FIB entries to a producer device and vice versa. In our NDN deployment, the producer initially creates all content objects. The 6LoWPAN producer configures a callback function as a CoAP endpoint for a particular URI to trigger a response.

b) *Multi-hop*: An extra forwarding device is added to the network topology, such that requests and responses between the consumer and producer traverse through the forwarder.

2) *Hardware & Software Platform*: We conduct all our experiments on typical class 2 [11] devices that feature an ARM Cortex-M0+ MCU with 32 kB RAM, 256 kB ROM and up to 48 MHz CPU frequency. Each device further provides an Atmel AT86RF233 [28] 2.4 GHz IEEE 802.15.4 radio transceiver. We set the radio transmission power to 0 dBm, the receiver sensitivity to -94 dBm and enable the *Smart Receiving* feature, an energy saving mode for idle listening. For our power consumption measurements we make use of on-board current measurement headers on each of our devices. We measure currents using a Keithley DMM7510 7½ digit graphical sampling multimeter with 1 MHz sampling rate and control it with external I/O lines to trigger start and stop from events generated by our network stack. Devices under test are powered by a regulated external DC power supply and connect via UART to a Linux control node to obtain experiment results.

In each experiment, our devices operate RIOT OS version 2018.10. Our NDN deployments use the CCN-lite [29] package and our 6LoWPAN experiments are based on the default GNRC network stack of RIOT OS. We integrate ICNLoWPAN into the 6LoWPAN module of RIOT OS to reuse the code base of the dispatching framework and link fragmentation.

3) *Name Configuration*: Name lengths proportionally affect processing times, packet lengths and consequently energy expenditure during transmissions. This especially impacts NDN as names are included in requests as well as in responses. We use two different names in our experiments to measure the effects of our stateless and stateful compression mechanisms.

a) *Name_{short}*: We use a short name with 4 components to denote temperature readings produced by a sensor. The name is of the form `/org/example/temp/idx`, where `idx` is an increasing number for each request. A CID is configured for `/org`, such that this prefix is elided from all messages.

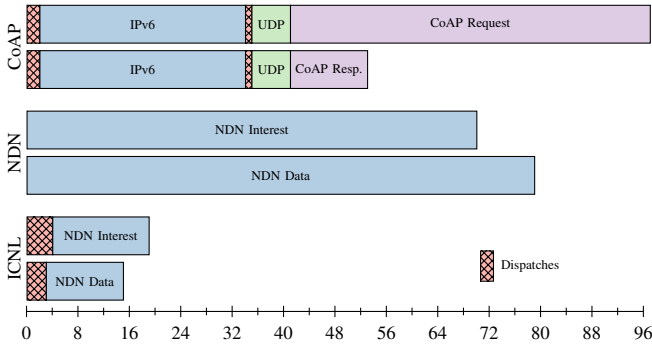


Fig. 8: Packet length and structure for different protocols.

b) *Name_{long}*: We use a long name with 10 components of the form */org/example/building/1/floor/4/room/481/temp/id_x*. A CID is configured for the prefix */org/example/building/1/floor/4/room/481* to elide a considerable portion of the name.

B. Theoretical Evaluation

In this first evaluation, we analyze NDN and CoAP packet sizes for a typical IoT scenario using *Name_{long}* as a CoAP endpoint and as a naming scheme for NDN. We further configure responses of both protocols to return 4-byte signed integer values that represent temperature sensor readings.

Fig. 8 depicts the actual packet sizes for each protocol. Our CoAP request has a packet size of 97 bytes, where 3 bytes are used for 6LoWPAN dispatches, 32 bytes for the compressed IPv6 header and 6 bytes for the UDP header. The remaining 56 bytes are used by the actual CoAP message. The respective CoAP response requires considerably less, which follows from CoAP omitting URIs in responses and using tokens to match against open requests on the requesting node. In our setup, CoAP uses 2-byte tokens for each request and returns the exact token in the response. The NDN Interest message nets to 70 bytes, whereas the returning data message requires 79 bytes. Contrary to CoAP, returning responses in NDN include the name of the request, or even a more specific and longer name. The displayed data message further contains empty Signature TLVs. The equivalent ICNLoWPAN compressed NDN messages are significantly shorter, where the Interest message reduces down to 19 bytes (72 % savings) and the data message down to 15 bytes (81 % savings). For Interests, this gain is mainly due to leveraging the configured stateful name compression and data messages naturally benefit from eliding the full returning name. In addition to the compressed messages, we require a 1-byte page dispatch, 1-byte ICNLoWPAN dispatch, and a 1-byte HopID for Interest and data. The compressed Interest packet further includes a 1-byte CID indicating the elided prefix for the stateful name compression.

C. Experimental Evaluation

The theoretical evaluation indicates that ICNLoWPAN substantially reduces packet sizes in Named-Data IoT networks

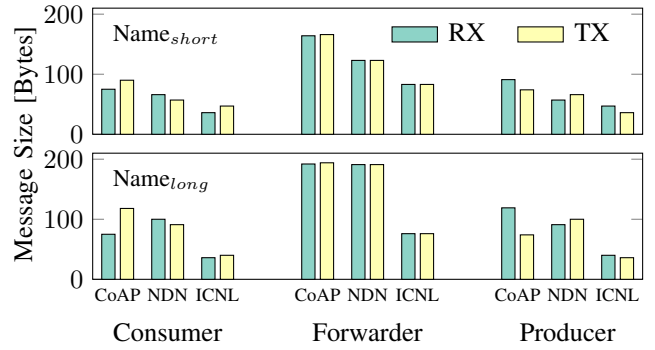


Fig. 9: Average bytes per request-response.

with pull-driven traffic patterns. In this experimental evaluation, we want to explore the effects of our convergence layer on resource-constrained nodes to gauge the feasibility for real-world low power networks. In particular, we measure (i) intra-stack processing times, (ii) average message overhead for a request-response handshake, (iii) energy expenditure during transmissions for single-hop and multi-hop deployments, (iv) and effects on reliability when periodically requesting sensor values while we enable randomized and bursty cross-traffic.

1) *Processing Times*: We first evaluate intra-stack processing times in a multi-hop deployment, where a consumer requests every 500 ms a specific temperature reading from a producer using *Name_{short}* and *Name_{long}*. Timestamps for the link layer, convergence layer (LoWPAN), and upper layers are recorded for each packet using a hardware timer on each device with μs precision. The link layer time depicts operations of the RIOT OS radio driver including SPI transfer of the packet to the radio frame buffer with 5 MHz. This measurement does not contain the actual transmission or reception of a frame over the wireless medium, as this procedure does not load the CPU. Time spent in the LoWPAN module includes the handling of dispatch headers and packet (de-)compression. Processing times for the network layer and beyond either include IPv6, UDP, and CoAP or NDN operations in addition to the actual application on top that issues or satisfies requests.

Fig. 7 displays experiment results for the different roles of a consumer, forwarder, and producer. We first observe that for the *Name_{short}* configuration, the additional processing overhead of ICNLoWPAN does not pay off on the link layer. Thus, some measurements with ICNLoWPAN take slightly more CPU resources than plain NDN. Conversely, savings on the link layer for our *Name_{long}* configuration visibly outperform the ICNLoWPAN processing overhead, especially for response packets by a decrease of $\approx 100 \mu\text{s}$ per packet.

2) *Message Sizes*: We now analyze the amount of bytes that are transmitted between consumer, forwarder, and producer when performing a request-response handshake in Fig. 9. The captured packets include the message lengths of Fig. 8 in addition to 21 bytes for an IEEE 802.15.4 header and 2 bytes for FCS. While comparing the *Name_{short}* and *Name_{long}* configurations, we observe an increase in sent bytes for the

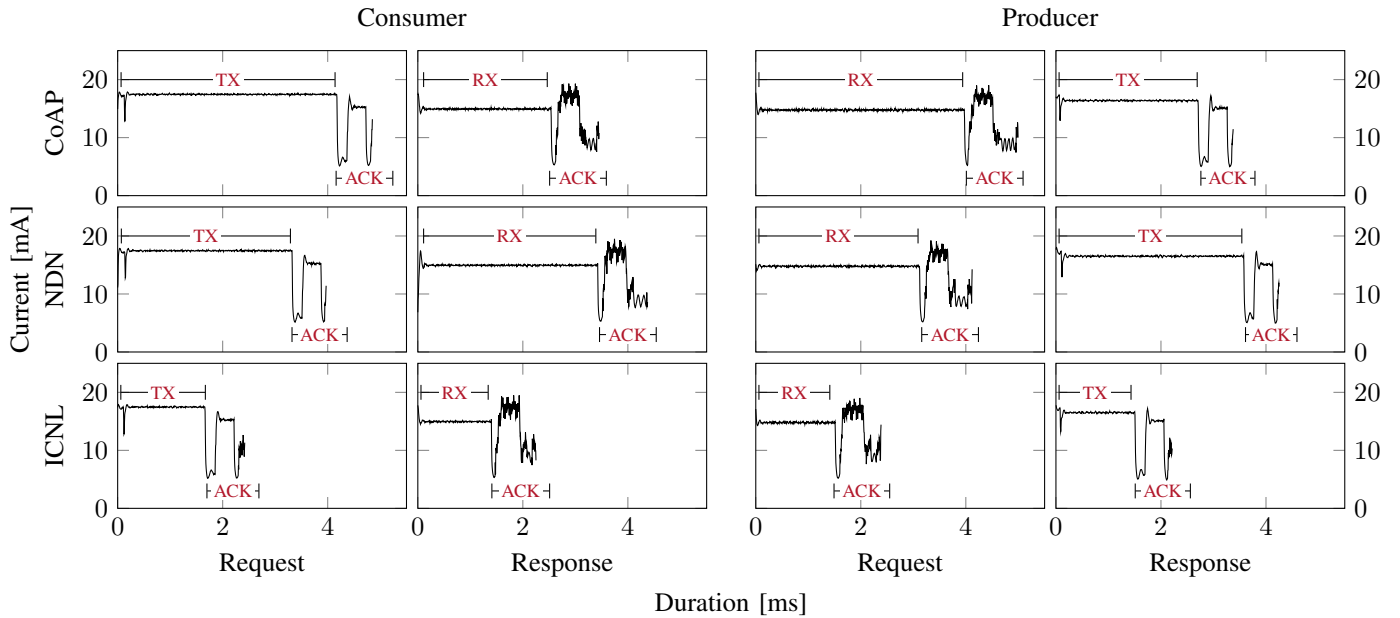


Fig. 10: Current consumption for send and receive operations.

	Consumer		Forwarder		Producer	
	$Name_{short}$	$Name_{long}$	$Name_{short}$	$Name_{long}$	$Name_{short}$	$Name_{long}$
CoAP	548.58 μ J	612.24 μ J	967.41 μ J	1072.07 μ J	464.73 μ J	517.96 μ J
NDN	526.23 μ J	687.26 μ J	880.68 μ J	1152.02 μ J	422.55 μ J	584.82 μ J
ICNL	466.09 μ J	487.32 μ J	769.17 μ J	773.97 μ J	369.84 μ J	395.19 μ J

TABLE I: Energy consumption in μ J

CoAP and NDN consumer, whereas ICNLoWPAN reduces the amount of sent bytes due to a better utilization of our name compression scheme. We interestingly see that the amount of sent bytes for the CoAP producer stagnates, while the amount for our NDN producer increases. This follows from the fact that CoAP responses do not include the URI component, but rather a fixed-length token to match against open requests. In contrast, the NDN data message does include the full name that is obtained from the request, which leads to significantly increased message sizes from $Name_{short}$ to $Name_{long}$.

3) *Energy Consumption*: ICNLoWPAN decreases the total amount of bytes over the air with both name configurations compared to NDN. We observe a strong reduction of bytes for responses at the producer for both name configurations, since the name is fully elided. The reduction is most prominent for the $Name_{long}$ configuration at the forwarder with a drop by 60 % from 191 bytes down to 76 bytes. This gain is to be expected, considering that the forwarder is involved in four transmissions.

The packet length during transmission has an immediate effect on the energy consumption. We measure the current while sending and receiving messages for each role separately in a single-hop setup and display the results in Fig. 10 for $Name_{long}$. The graphs involve transmission over the

wireless, radio turnaround time as well as link layer frame acknowledgment. In our setup, sending draws slightly higher current than receiving and the duration of each transmission depends on the packet length. In fact, the duration of each depicted measurement correlates with the respective message size displayed in Fig. 8 and the results showed in Fig. 9, so that larger messages yield longer periods of operation for sending and receiving.

On the consumer, we observe that our CoAP request requires 5 ms to complete, while the respective NDN request is transmitted in 4 ms, including the reception of acknowledgments for both. Conversely, the CoAP response is received by the consumer in 3.8 ms, while the NDN response completes in 4.2 ms, including the sending of acknowledgments. With ICNLoWPAN in operation, we notice a decrease of transmission times by around ≈ 50 % on the consumer due to compressed messages and the resulting shortened media utilization. As expected, the reduction for responses is more prominent due to fully eliding Name TLVs. On the producer, we naturally observe mirrored results for each operation.

Given the fact that current draws for transmissions in Fig. 10 are mainly similar, the actual energy consumption is predominated by the transmission durations. We thus analyze the overall power consumption of a request-response handshake

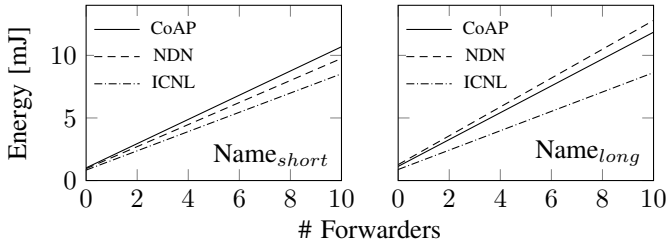


Fig. 11: Total energy consumption for multi-hop networks.

in a multi-hop setup for $Name_{short}$ and $Name_{long}$ including processing times and transmission durations in Tab. I. Our results indicate an increase in energy usage for each role with the $Name_{long}$ configuration compared to the $Name_{short}$ configuration. We further notice that our producer spends the least amount of energy, followed by the consumer, and our forwarder expends nearly double the amount of energy than the producer. The increased power consumption is inherently consistent with the fact that the forwarder is involved in two request and subsequently in two response transmissions.

The NDN consumer device uses 4 % less energy for $Name_{short}$ and 12 % more energy for $Name_{long}$ compared to the CoAP consumer. This turnaround in energy expenditure for $Name_{long}$ is twofold. 1) NDN has a more verbose name encoding than CoAP and 2) CoAP does not include the URI in the response. ICNLoWPAN reduces energy usage of our NDN consumer by 11 % for $Name_{short}$ and 29 % for $Name_{long}$. Our NDN producer consumes 9 % less energy for $Name_{short}$ and 13 % more energy for $Name_{long}$ compared to our CoAP producer. The energy consumption reduces by 12 % for $Name_{short}$ and 32 % for $Name_{long}$ with an enabled ICNLoWPAN operation compared to NDN. Since our forwarder interacts with four transmissions, we observe a natural increase in overall expenditures. The NDN forwarder consumes 9 % less power for $Name_{short}$ and 7 % more energy for $Name_{long}$ compared to the CoAP forwarder. In contrast, ICNLoWPAN reduces the expenditure by 13 % for $Name_{short}$ and 33 % for $Name_{long}$. The trend that ICNLoWPAN yields higher energy savings for $Name_{long}$ becomes apparent.

Finally, we calculate the energy consumption of a full request-response handshake for a multi-hop setup with a varying number of forwarders between a consumer and producer in Fig. 11. For all setups, we see an increase in expenditure for CoAP, NDN, and ICNLoWPAN with an increasing number of forwarders. We again notice that the power consumption for NDN surpasses the consumption for CoAP using the $Name_{long}$ configuration due to returning Name TLVs in the response. ICNLoWPAN clearly reduces the overall energy consumption of an NDN handshake for both configurations $Name_{short}$ and $Name_{long}$. Interestingly, despite the increase in power consumption for NDN versus CoAP, ICNLoWPAN manages to cut expenditures by ≈ 25 % for a setup with ten forwarders compared to NDN.

4) *Reliability*: In this experiment we investigate the reliability of a typical data retrieval setup in our single-hop

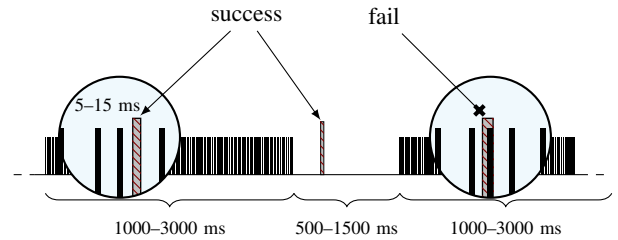


Fig. 12: Radio interference due to bursty cross-traffic.

deployment, where a consumer periodically requests temperature values from a producer every 300 ms. We additionally generate bursty cross-traffic in randomized intervals to a third party in order to mimic a dense network with multiple devices that periodically request sensor readings. Fig. 12 illustrates the configured traffic pattern of our cross-traffic with each burst consisting of 200 UDP packets in succession with a 5–15 ms delay in between each transmission and a radio silence interval of 500–1500 ms. It is worth noting that our cross-traffic only adds wireless interference as the IEEE 802.15.4 frames are not delivered to the devices due to automatic MAC address filtering in hardware of the radio module. We further disable CSMA/CA and frame retransmissions of the radio transceiver to explore reliability gains without distortions of automatic corrective actions in hardware. Indeed, disabling such hardware features in real deployments is not a far-fetched scenario as sophisticated link layer protocols that employ time-slotted algorithms must satisfy strict time constraints, while CSMA/CA and retransmissions lead to indeterministic and drifting media accesses. In fact, several radio transceivers that dual operate IEEE 802.15.4 and Bluetooth Low Energy do not even support such mechanisms in hardware, but rather rely on software implementations.

Tab. II lists the packet reception ratio (PRR) for consumer and producer roles using the $Name_{long}$ configuration. For each deployment the number of received requests on the producer lies within ≈ 93 –94 %. Conversely, the percentage of successfully received responses on each consumer clearly varies between ≈ 73 –76 % for CoAP and NDN and ≈ 93 % for our ICNLoWPAN operation.

The performance gain of ICNLoWPAN results from strongly compressed packets which lead to a significantly reduced on-air time for the low power wireless transmission (see Fig. 10). This reduces collision probability with interferer traffic, especially when responses are sent during a burst as shown in Fig. 12. Furthermore, responses of a producer always follow the successful reception of a request. Reduced transaction times with ICNLoWPAN leave more time to the

	CoAP	NDN	ICNL
Producer	93.53 %	94.05 %	94.40 %
Consumer	73.25 %	75.98 %	93.06 %

TABLE II: Packet Reception Ratio (PRR)

next interferer transmission within a burst, which further reduces the probability of overlapping cross-traffic compared to the NDN and CoAP operation.

V. CONCLUSIONS

IoT networking has proven to benefit from information-centric architectures in several directions. In this paper, we worked out the components for adapting NDN to a LoWPAN edge: compression, framing, and fragmentation. By leveraging the NDN stateful forwarding for compression on path, we could again take specific advantage of the information-centric approach in low power lossy IoT networks.

Extensive measurements of our implementation on IoT hardware revealed these benefits in comparison with plain NDN and the IP world (CoAP over 6LoWPAN). Our experimental results clearly showed that ICNLoWPAN outperforms NDN and CoAP in terms of media utilization as well as energy consumption. ICNLoWPAN further reduces end-to-end latencies in multi-hop scenarios, and contributes to an improved reliability in lossy environments while preserving battery resources. Depending on the use case, savings typically range from 20 % to 33 %.

With these results, we hope to contribute insights to the community and to encourage deployment of NDN in the constrained IoT. Our future work will concentrate on extending ICNLoWPAN to different low power link technologies, such as BLE and align with adaptation layers for wide-area and cellular technologies to enable LoRA and NB-IoT for a versatile, efficient, and robust information-centric Internet of Things.

ACKNOWLEDGMENTS

This work was supported in parts by the German Federal Ministry of Research and Education within the project *I3: Information-centric Networking for the Industrial Internet*.

REFERENCES

- [1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A Survey of Information-Centric Networking," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26–36, July 2012.
- [2] 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 Information-Centric Networking (ICN-2014)*. N.Y.: ACM, Sept. 2014, pp. 77–86.
- [3] W. Shang, A. Bannis, T. Liang, Z. Wang, Y. Yu, A. Afanasyev, J. Thompson, J. Burke, B. Zhang, and L. Zhang, "Named Data Networking of Things (Invited Paper)," in *Proc. of IEEE International Conf. on Internet-of-Things Design and Implementation (IoTDI)*. Los Alamitos, CA, USA: IEEE Computer Society, 2016, pp. 117–128.
- [4] E. M. Schooler, D. Zage, J. Sedayao, H. Moustafa, A. Brown, and M. Ambrosin, "An Architectural Vision for a Data-Centric IoT: Rethinking Things, Trust and Clouds," in *IEEE 37th Intern. Conference on Distributed Computing Systems (ICDCS)*. Piscataway, NJ, USA: IEEE, June 2017, pp. 1717–1728.
- [5] K. Pentikousis, B. Ohlman, D. Corujo, G. Boggia, G. Tyson, E. Davies, A. Molinaro, and S. Eum, "Information-Centric Networking: Baseline Scenarios," IETF, RFC 7476, March 2015.
- [6] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, "Networking Named Content," in *5th Int. Conf. on emerging Networking Experiments and Technologies (ACM CoNEXT'09)*. New York, NY, USA: ACM, Dec. 2009, pp. 1–12.

- [7] C. Gündogan, P. Kietzmann, M. Lenders, H. Petersen, T. C. Schmidt, and M. Wählisch, "NDN, CoAP, and MQTT: A Comparative Measurement Study in the IoT," in *Proc. of 5th ACM Conference on Information-Centric Networking (ICN)*. N.Y.: ACM, Sept. 2018.
- [8] Z. Shelby, K. Hartke, and C. Bormann, "The Constrained Application Protocol (CoAP)," IETF, RFC 7252, June 2014.
- [9] A. Stanford-Clark and H. L. Truong, "MQTT For Sensor Networks (MQTT-SN) Version 1.2," IBM, Protocol Specification, November 2013. [Online]. Available: http://mqtt.org/new/wp-content/uploads/2009/06/MQTT-SN_spec_v1.2.pdf
- [10] D. Kutscher, et al., "Information-Centric Networking (ICN) Research Challenges," IETF, RFC 7927, July 2016.
- [11] C. Bormann, M. Ersue, and A. Keranen, "Terminology for Constrained-Node Networks," IETF, RFC 7228, May 2014.
- [12] IEEE 802.15 Working Group, "IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)," IEEE, New York, NY, USA, Tech. Rep. IEEE Std 802.15.4™–2011, Sep 2011.
- [13] Bluetooth Special Interest Group, "Bluetooth Core Specification Version 5.1," Tech. Rep., January 2019. [Online]. Available: <https://www.bluetooth.com/specifications/bluetooth-core-specification>
- [14] LoRa Alliance, "A Technical Overview of LoRa and LoRaWAN," Tech. Rep., November 2015. [Online]. Available: <https://loro-alliance.org/sites/default/files/2018-04/what-is-lorawan.pdf>
- [15] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," IETF, RFC 4944, September 2007.
- [16] J. Hui and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks," IETF, RFC 6282, September 2011.
- [17] E. Baccelli, O. Hahm, M. Günes, M. Wählisch, and T. C. Schmidt, "RIOT OS: Towards an OS for the Internet of Things," in *Proc. of the 32nd IEEE INFOCOM. Poster*. IEEE Press, 2013, pp. 79–80.
- [18] E. Baccelli, C. Gündogan, O. Hahm, P. Kietzmann, M. Lenders, H. Petersen, K. Schleiser, T. C. Schmidt, and M. Wählisch, "RIOT: an Open Source Operating System for Low-end Embedded Devices in the IoT," *IEEE Internet of Things Journal*, vol. 5, no. 6, pp. 4428–4440, December 2018.
- [19] M. Gritter and D. R. Cheriton, "An Architecture for Content Routing Support in the Internet," in *Proc. USITS'01*. Berkeley, CA, USA: USENIX Association, 2001, pp. 4–4.
- [20] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A Data-Oriented (and beyond) Network Architecture," *SIGCOMM Computer Communications Review*, vol. 37, no. 4, pp. 181–192, 2007.
- [21] C. Bormann, "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)," IETF, RFC 7400, November 2014.
- [22] W. Shang, A. Afanasyev, and L. Zhang, "The Design and Implementation of the NDN Protocol Stack for RIOT-OS," in *Proc. of IEEE GLOBECOM 2016*. Washington, DC, USA: IEEE, 2016, pp. 1–6.
- [23] M. Mosko and C. Tschudin, "ICN 'Begin-End' Hop by Hop Fragmentation," IETF, Internet-Draft – work in progress 02, December 2016.
- [24] J. Shi and B. Zhang, "NDNLP: A Link Protocol for NDN," NDN Team, NDN, Technical Report NDN-0006, July 2012.
- [25] C. Ghali, A. Narayanan, D. Oran, G. Tsudik, and C. A. Wood, "Secure Fragmentation for Content-Centric Networks," in *Proc. of the 14th International Symposium on Network Computing and Applications*, 2015, pp. 47–56.
- [26] Y. Yang and T. Song, "Local Name Translation for Succinct Communication Towards Named Data Networking of Things," *IEEE Communications Letters*, vol. 22, pp. 2551–2554, Dec. 2018.
- [27] P. Thubert and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch," IETF, RFC 8025, November 2016.
- [28] Atmel, *Low Power 2.4 GHz Transceiver for ZigBee, IEEE 802.15.4, 6LoWPAN, RF4CE, SP100, WirelessHART, and ISM Applications*, Atmel Corporation, September 2009. [Online]. Available: <http://www.atmel.com/images/doc81111.pdf>
- [29] C. Tschudin, C. Scherb et al., "CCN Lite: Lightweight implementation of the Content Centric Networking protocol," 2018. [Online]. Available: <http://ccn-lite.net>