Rethinking LoRa for the IoT: An InformationCentric Approach

Peter Kietzmann, José Alamos, Dirk Kutscher, Thomas C. Schmidt, and Matthias Wählisch

The authors present LoRa-ICN, a new long-range communication system that provides a versatile data-oriented integration of battery-driven LoRa nodes into the Internet of Things.

ABSTRACT

In this article, we present LoRa-ICN, a new long-range communication system that provides a versatile data-oriented integration of battery-driven LoRa nodes into the Internet of Things (IoT). LoRa-ICN builds on two paradigms: information-centric networking (ICN), which enables more direct, data-oriented communication between Internet systems and the low-power wireless domain, and 802.15.4 DSME, which is an IoT MAC layer that facilitates reliable LoRa transmissions. While the combination of LoRa and DSME is generally better suited for bi-directional end-to-end communication, it still incurs considerable long and variable transmission latencies, and challenges the network layer transition between the power-constrained wireless domain and the Internet. Our design and implementation on actual off-the-shelf IoT hardware includes extensions to ICN that enable delay-tolerant data retrieval between LoRa nodes and an application on the Internet. An experimental comparison between default ICN mechanisms and our extensions shows that LoRa-ICN is able to achieve a high data delivery rate, while dealing with the higher latencies explicitly, thus providing a viable option for re-imaging LoRa networks with a data-oriented, Internet-friendly approach.

INTRODUCTION

The Internet of Things (IoT) interconnects numerous sensors and actuators either locally or across the global Internet. From an application perspective, IoT systems are inherently data-oriented, that is, their purpose is often to provide access to named sensor data and control interfaces. From a device and communication perspective, things in the IoT are resource-constrained devices that are commonly powered by a small battery and communicate wirelessly.

LoRaWAN systems today integrate the LoRa physical layer with the LoRaWAN MAC layer and corresponding infrastructure support. Among the IoT radio technologies, LoRa is a versatile and popular candidate [1, 2] since it provides a physical layer that allows for data transmission over multiple kilometers with minimal energy consumption. At the same time, the high LoRa receiver sensitivity enables packet reception in noisy environments, which makes it attractive for industrial deployments. On the downside, LoRa achieves only low data rates requiring long on-air times, and significantly higher latencies compared to radios that are typically used for Internet access.

The LoRaWAN MAC layer and network architecture that is often used in LoRa deployments, thus provide a vertically integrated sensor data delivery service on top of the LoRa PHY that implements media access and end-to-end network connectivity. Unfortunately, LoRaWAN cannot utilize the LoRa PHY to its best potential with respect to throughput and robustness and is mostly used for upstream-only communication. It is not intended to directly interconnect with the Internet, but relies on a bespoke middlebox architecture consisting of gateways and *network servers*. Overall LoRaWAN has the following main problems, as depicted in Fig. 1a.

Centralization around a network server prevents data sharing between users, across distributed applications, and requires permanent infrastructure backhaul of the wireless access network. Uplink-oriented and uncoordinated communication leads to wireless interference. Downlink traffic is rarely available in practice and suffers from scalability issues.

This article presents an overview about recent advancements to enable data-centric, long-range IoT communication based on LoRa. The proposed network system aims for delay-tolerant, bi-directional communication in the presence of vastly longer latencies and lower bandwidth compared to regular Internet systems — without relying on vertically integrated middlebox-based architectures (Fig. 1b). The resulting system resolves current LoRaWAN performance issues using two main building blocks: a new network layer based on Information-centric Networking (ICN) and a new MAC layer.

Originally designed for non-constrained wired networks to abandon the end-to-end paradigm and access data only by names instead of IP endpoints, ICN migrated to the constrained wireless IoT over the past years. ICN still lacks a lower layer definition but provides mechanisms that are beneficial for the challenging LoRa domain: Decoupling of content from endpoints separates data access from physical infrastructure; Inherent content caching and replication potentially reduce link load, thus, wireless interference, and it preserves battery resources. The ICN-LoRa system presented in this article bases its design on IEEE 802.15.4 DSME [3], which was originally designed for low-power personal area networks. This MAC handles media access reliably using time- and frequency multiplexing, and enables reliable bi-directional communication.

ISSN: 0163-6804
Digital Object Identifier: 10.1109/MCOM.001.2300379

Peter Kietzmann, José Alamos, and Thomas C. Schmidt are with HAW Hamburg, Germany; Dirk Kutscher is with HKUST(Guangzhou), China; Matthias Wählisch is with TU Dresden and Barkhausen Institut, Germany.

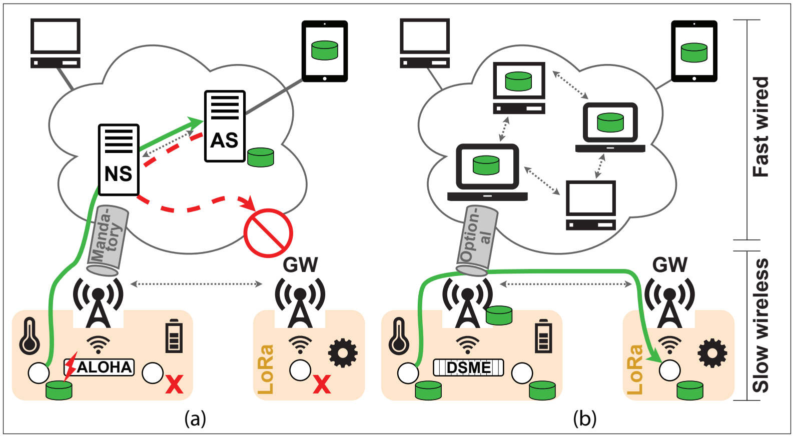

FIGURE 1. a) server-centric (LoRaWAN); b) content-centric (LoRa-ICN) LoRa deployments. AS: Application server, NS: Network server, GW: Gateway.

Synergizing the advantages of LoRa, DSME, and ICN enables delay-tolerant, bi-directional LoRa communication, wich enhances many existing IoT applications. Wide area data retrieval and control as for solar power stations or smart street lighting systems are facilitated by the new MAC and its ICN integration. High voltage overhead line monitoring connecting voltage sensors and transformers relies on high data reliability, even under intermittent connectivity or loss. ICN achieves this, employing content caching and replication. Traveling container monitoring [4] is challenging due to mobility and interference from metallic surfaces, where LoRa surpasses other radio systems. Decoupling content from its location for mobile containers and an adaptation to long producer delays are naturally contributed by LoRa-ICN.

In the remainder of this article, we provide the essential technical background and challenges to design a LoRa-ICN system. We identify the key performance potentials of five protocol variants based on an implementation in RIOT OS [5] and experiments on off-the-shelf IoT devices. Finally, we conclude and discuss directions for future research.

TECHNICAL BACKGROUND AND IOT CHALLENGES

In this section, we present background about LoRa, the 802.14.5 DSME MAC layer, and Information-Centric Networking — all three components that form the base of LoRa-ICN.

LONG RANGE RADIO SYSTEM

LoRa defines a chirp spread spectrum modulation that enables a long transmission range (kilometers) at low energy consumption (millijoules), but with very limited throughput (bits per second). Spread spectrum techniques are robust against wireless disturbances such as multipath fading, or doppler effect. This is particularly useful for operation in the unlicensed sub-GHz ISM band in which LoRa operates, causing background noise from parallel transmissions. In addition to cross traffic, industrial deployments introduce electromagnetic interference.

The LoRa modulation is highly configurable: the spreading factor (SF), code rate, and bandwidth can be adjusted and directly affect the time on air and data rate. Varying the center-frequencies in the sub-GHz ISM band constrains the duty cycle to 0.1–10%, which may limit the maximum effective bitrate of the physical layer to only 0.25 bit/s.

LoRaWAN defines three LoRa MAC operating modes. Class A is purely uplink oriented, allowing for maximum device sleep. Nodes send uplink packets using the best-effort ALOHA protocol and can receive downlink traffic within two subsequent slots. Class B adds periodic downlink slots at low energy, but exhibits scalability, and reliability issues [6, 7]. Class C is unsuitable for battery-driven devices. LoRaWAN relies on network servers on the Internet that schedule and terminate LoRaWAN protocol sessions. LoRaWAN gateways are peripherals of such *network servers* and relay traffic between radio networks and the Internet.

Challenges: The current LoRaWAN system design, based on the wide LoRa transmission range and long on air times, combined with best-effort media access, can induce wireless interference and transmission errors. The limited downlink capabilities prevent many IoT use cases for energy-constrained nodes. Centralization around a network server complicates data sharing at the edge and the vertically integrated architecture prevents direct communication with data consumers or device controllers in the same LoRa network, and over the Internet.

LOW POWER MAC LAYER BASED ON IEEE 802.14.5 DSME

An 802.14.5 DSME coordinator starts network formation and emits beacons that initiate a time synchronized slot frame structure, enabling time division and frequency multiplex for nodes that joined the network. A 'superframe' consists of three transmission periods:

- A beacon period
- A contention-access period that enables sporadic communication on a common frequency channel
- A contention-free period, which consists of guaranteed time slots, multiplexed across orthogonal radio channels.

The transmission slots have to be negotiated with the coordinator before use. Time division enables device duty cycling to save energy, and exclusive transmission slots provide deterministic latency with high reliability. DSME supports peer-to-peer and cluster tree topologies. A network can be extended by adding coordinators that share the same area.

Challenges: The MAC builds on top of 802.15.4 PHY properties (e.g., symbol time) and requires an 802.15.4-compliant adaptation layer to operate on different radio technologies. LoRa is highly configurable and exhibits significantly different transmission properties. The MAC mapping needs to define compatible LoRa PHY settings to maintain a synchronized slot frame structure across many nodes that join the same DSME network. Furthermore, due to long LoRa times on-air, the duration of a MAC frame notably increases, that is, the latency introduced by time-slotting. This challenges higher-layer protocols, which need to cope with these delays.

INFORMATION-CENTRIC NETWORKING

Information-Centric Networking (ICN) is a network architecture providing access to named (and authenticated) data in the network. This concept has been implemented by protocols such as Named Data Networking (NDN) [8] (Fig. 2a). In NDN, two message types, Interest and Data, realize a request-response pattern, and three data structures act as the basis for implementing the

To enable resource-efficient LoRa-ICN edge communication, we define new node- and gateway behavior and extend the basic ICN protocol flow to handle long and vastly different producer delays, mainly introduced by slow LoRa transmissions and time-slotting of DSME.

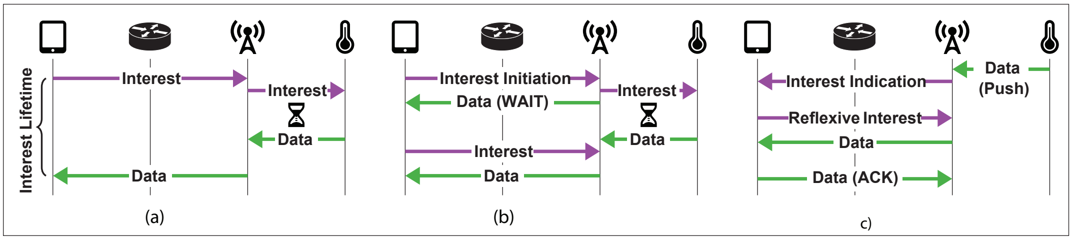

FIGURE 2. Sequence flows of a default (Vanilla) ICN request using NDN, and two extensions for delay-tolerant data retrieval in LoRa-ICN: a) Vanilla ICN Data Retrieval; b) Delay-tolerant Data Retrieval; c) Reflexive Push.

protocol logic: Forwarding Information Base (FIB), Pending Interest Table (PIT), and Content Store (CS). A consumer sends an Interest for a named data object. Nodes that receive an Interest first perform a name lookup in their CS and directly respond with data, if the data object is available. Otherwise, the Interest is forwarded according to the FIB and creates a new PIT entry. If a PIT entry for the Interest name already exists, forwarders can decide to add the new information to the existing PIT entry and not forward the Interest again (*Interest aggregation*). A distant producer serves the Interest with a data packet that contains the name, the data, and a signature. Data follows the reverse routing path noted in the PIT entries. PIT entries are removed while data is forwarded to the consumer. Data packets without corresponding PIT entry are dropped.

Reliability is achieved through Interest retransmissions that can be initiated by consumers or forwarders in the network. Typically, retransmissions are timeout-triggered, that is, when data is not received within a specified time interval. This mechanism in conjunction with the Interest aggregation option implies that PIT state must eventually expire (to ensure resynchronization of the forwarding plane). The Interest lifetime can be specified in an `InterestLifetime` field in NDN (the default is 4 s).

Challenges: Long producer delays as introduced by LoRa, require long-lived PIT state to forward the data packet to the consumer. Regular forwarders might object to save many PIT entries for an untypically long time (more than 4 seconds). This can lead to losses owing to the expired forwarding state. Interest retransmissions are aggregated as long as the PIT entry remains active, complicating recovery after real packet loss (on the wire). With ICN-idiomatic behavior, a periodic consumer that faces a slow sporadic consumer would have to poll the network, which conflicts with the limited LoRa radio resources. Hence, the value of the `InterestLifetime` and retransmission timeout are crucial in delay-tolerant ICN networks. Estimating suitable poll intervals and timeout values would have to be adjusted constantly, introducing overhead to the consumer application and the network.

INTEGRATING LORA WITH ICN

LoRa-ICN aims to overcome the limitations of LoRaWAN. In our new IoT system architecture [9], an application on the Internet can directly access data of constrained LoRa nodes, without requiring changes to the ICN Internet itself or application layer relays. Each LoRa access network is served by one gateway, which acts as an application-agnostic caching forwarder and DSME coordinator. Maintaining the slotframe structure introduces an overhead to LoRa nodes, increasing the energy consumption by a factor of approximately three [10] compared to LoRaWAN class A. This is still feasible for battery operated nodes, though, and the DSME MAC schedule allows leveraging the power conservation features of the PHY, in contrast to LoRaWAN class C. A gateway connects the wireless network to the ‘regular’ ICN Internet. The constrained nodes implement the corresponding ICN producer/consumer logic. Connecting a constrained narrowband network to an Internet face is particularly challenging, which motivates us to re-visit the overall system architecture. We first introduce cross-layer mappings that enable ICN message transport over DSME, and its reliable energy-conserving transmission via LoRa radios. To enable resource-efficient LoRa-ICN edge communication, we define new node- and gateway behavior and extend the basic ICN protocol flow to handle long and vastly different producer delays, mainly introduced by slow LoRa transmissions and time-slotting of DSME. In terms of energy conservation, our LoRa-ICN devices achieve a nodal lifetime of approximately 230 — 384 days when using an AA alkaline battery with 2800 mAh [9].

MAPPING BETWEEN LAYERS

DSME-LoRa: Mapping the DSME MAC to the LoRa PHY requires three adaptations [11].

1. We assign 16 frequencies in the EU868 region to LoRa channels, with a 200 kHz channel spacing. One channel is located in a band with 10% duty cycle restriction and used for periodic beacon transmission by the coordinator, contention-access-, and contention-free traffic. The other 15 channels exhibit a 1% restriction and exclusively serve contention-free media access. All channels utilize SF7, 125 kHz bandwidth, and code rate 4/5, which results in a PHY bit rate of ≈ 5.5 kb/s, and a symbol time of ≈ 1 ms. Multiple PHY settings are configurable through channel pages, which affects the timing behavior (see 2) and requires more dynamic changes to the slotframe instantiation and scheduling logic. This bears the potential to further improve the transmission range, total throughput, and error resilience (future work).
2. The frame mapping configures the DSME symbol time to align with LoRa (1 ms, see 1). In order to transmit full 127 Bytes 802.15.4 frames, we set the superframe order to 3, which results in a duration of 7.68 s. With four superframes per multi-superframe, the slotframe structure repeats every 30.72 s and provides 28 timeslots · 16 frequency channels = 448 exclusive transmission cells. DSME is highly configurable, and the coordinator can dictate changes in the slotframe structure, increasing the number of transmission slots or decreasing the latency. Hence,

the expected DSME-LoRa round-trip times can vary significantly.

3. DSME uses clear channel assessment for contention-based media access. Our adaptation layer maps to the LoRa channel activity detection feature, which perceives the presence of a LoRa preamble on the air; however, it only detects the preamble and not an ongoing frame transmission. This still improves the reception rate by more than 25% when compared to the ALOHA-based media access in LoRaWAN [11].

ICN-DSME: DSME provides three options for transporting ICN messages: contention access, exclusive time slots, or beacons. The contention-access period is prone to collisions; hence, we only utilize it for registration purposes (below) and negotiation of contention-free transmission slots as per regular MAC behavior (see [3], Section 6.11.5). Later, we set up a static assignment for keeping our experiments reproducible.

Bi-directional communications, that is, ICN request-response, utilize two contention-free transmission slots. An Interest slot is followed by a data slot within the same superframe, keeping the round-trip time possibly low when data is available. Otherwise, the data response would increase the round trip time by at least one slotframe duration, further challenging the delay-unaware ICN protocol variants. Unidirectional uplink traffic, that is, unsolicited data push (below), requires only a single transmission slot per node. This doubles the number of available slots in a superframe compared with the request-response paradigm, and saves scarce link and energy resources. Downstream multicast traffic, for example, for distributing firmware updates to LoRa nodes, is implemented resource-efficiently by overloading beacons. Data broadcast can satisfy many pending Interests even if aggregated during a beacon period with a single downlink data packet, without requiring additional resources. Note that beacons are always transmitted to maintain time synchronization.

ICN Extensions

Gateway and Node Behavior: A LoRa node synchronizes with the slotframe structure of the MAC, advertised by the coordinator, that is, the gateway. After joining the DSME network, transmission slots can be negotiated, if not assigned statically. LoRa nodes register a node prefix at the gateway using ICN Interest/data (ACK), adopting NDN prefix registration [12]. They can only produce and serve content under that prefix. The gateway establishes a downlink routing entry in the FIB and acts as an ICN forwarder and LoRa node custodian with content cache. Gateways merely forward Interests with registered LoRa prefixes to the wireless domain. Forwarding and caching reduce (re-)transmissions of Interest and data messages, as in other ICN scenarios. This reduces wireless link load and packet processing on constrained nodes. Additionally, the gateway leverages knowledge about the long-delay domain of DSME-LoRa and lifts the PIT expiration time accordingly, to prevent premature state expiration. Duplicate Interests with the same name are aggregated (i.e., suppressed) as long as the PIT state remains active. Interests with unknown prefixes are naturally not forwarded without a missing FIB entry, thereby saving wireless link resources.

Consumer-Initiated Data Retrieval: Long and differing producer round trips challenge the ICN data retrieval. A node on a "regular" ICN Internet may remove PIT state prematurely, and fall back to data polling. To enable effective end-to-end ICN communication between an unconstrained Internet consumer and a slow, constrained LoRa node, our *delay-tolerant retrieval* extension leverages the concept of *RICE* [13], which supports long data retrieval delays (Fig. 2b). Thereby, our gateway acts as a proxy for long-delay producers and returns a distinct retry-time instruction (WAIT) on the first request of a content item, using its coordinator knowledge of the DSME state, that is, the time until the next transmission slot of a node is scheduled. In parallel, the request is forwarded to the (registered) LoRa node, and the returned data is cached. The estimated retrieval time in the first gateway response enables the consumer application to set an appropriate retry timer, without the need for specific producer knowledge or variably long delays. A distinct Interest retransmission at the given time is served from the gateway cache with high probability.

Producer-Initiated Data Retrieval: IoT nodes must save energy and produce data only occasionally, following their sleep cycles. This challenges ICN-idiomatic request-response communication. Our *Reflexive push* extension (Fig. 2c) enables a local unsolicited data push from registered LoRa nodes to the gateway, halving the number of resource-intensive wireless transmissions on the LoRa link. Constrained nodes wake up only on data availability, which reduces power consumption. Malicious publishers can simply be excluded by the coordinator by unregistering the node and muting its transmission slot. The gateway caches LoRa data in its content store according to caching rules. After placing content in the cache, the procedure leverages the *reflexive forwarding* extension to ICN [13]. It follows three steps:

1. The gateway forwards the data name to a pre-configured node on the Internet using the semantics of an Interest packet. Forwarders of that Interest note a temporary downstream routing entry in their FIB.
2. The consumer application on the Internet returns a reflexive Interest with that name, utilizing the forwarding information in the FIB, to retrieve the announced IoT data object. The request is served by the gateway, which results in an ICN-idiomatic protocol flow. Additionally, the temporary FIB state on forwarders is removed.
3. An optional data ACK terminates the initial indicating Interest.

Key Performance Results

Experiment Setup

We deployed a long-range producer application on off-the-shelf IoT hardware, that is, a Nordic nRF52840 microcontroller with a Semtech SX 1276 LoRa transceiver. The sensor node is operated by RIOT OS and our network stack (Fig. 3), consisting of the ICN implementation CCN-lite with our extensions, and the openDSME MAC, which we configure in agreement with above. Hence, the slotframe duration is ≈ 30 s, determining the upper bound of the inter-packet interval of a node. Our gateway acts

Bi-directional communications, that is, ICN request-response, utilize two contention-free transmission slots. An Interest slot is followed by a data slot within the same superframe, keeping the round-trip time possibly low when data is available.

This extension enables arbitrary producer delays without the need for long-lived forwarding state.

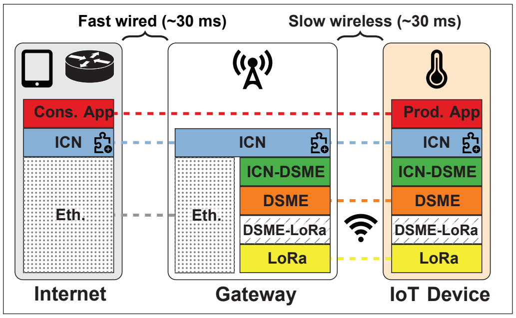

FIGURE 3. LoRa-ICN stacks on different devices with varying resources and network latencies.

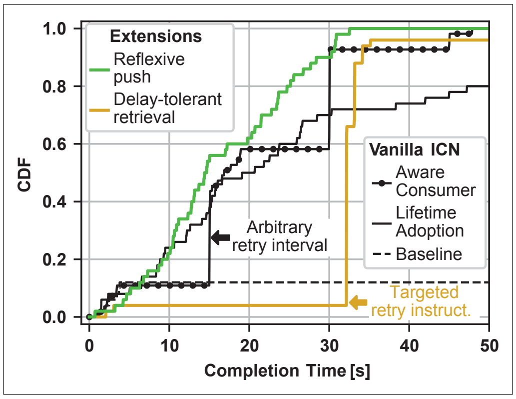

FIGURE 4. Time to content arrival with long producer delays. Vanilla ICN in varying configurations and our extensions.

as a coordinator for wireless nodes and connects them to a Linux-based workstation that bridges to a virtual TAP bridge via Ethernet. The fast Internet is emulated by two nodes in *Mininet*, one forwarder, and a consumer with an IoT application, exposing a latency of ≈ 30 ms and a loss of 5% per link.

Our measurement applications use unique content names. Data contains either a varying integer sensor reading or an ACK, NACK, or WAIT retry instruction. Every content item is transmitted once, with an average interval of one minute. We jitter transmissions (uniformly at 60 ± 10 s), to hit the DSME slotframe at different states.

Protocol Configuration

We compare Vanilla ICN protocol mechanisms in three configurations with the two delay-tolerant extensions, introduced earlier, and thereby focus on the typical IoT use case of data retrieval from a constrained LoRa node to an application on the "regular" ICN Internet.

Baseline acts as a default scenario with common parameter settings. An `InterestLifetime` of 4 s sets the PIT timeout on forwarders and the consumer retransmits Interests in a 1 s interval.

Aware Consumer depicts the case of an IoT application that is aware of the long producer delay and adjusts the Interest retransmission interval to 15 s. Forwarders stick to the standard PIT timeout value of 4 s as indicated by the `InterestLifetime` field.

Lifetime Adoption behaves similarly to the *Aware Consumer*, however, this time forwarders apply a long `InterestLifetime` of 60 s to set their PIT timer.

Delay-tolerant retrieval is based on a retry instruction by the gateway. In this case, neither polling nor estimating round trips at the consumer are used. Forwarders do not set an untypically long PIT timeout value.

Reflexive push inverts the transaction flow. Only on available IoT data in the gateway cache, the data announcement triggers a reflexive Interest request on the consumer, which can be satisfied immediately, without increasing timeouts nor poll intervals.

Evaluation Discussion

Figure 4 displays the completion times of successful transactions. Data losses result in infinite completion times and are inherently reflected by the gap between the end value of each graph and 1.0 (100%).

Vanilla ICN: The *Baseline* scenario exhibits a loss ratio greater than 85%. The expired PIT state on forwarders prevents data forwarding, and retransmissions on a shorter timescale than the DSME-LoRa round trip time do not significantly improve the success ratio. These losses can be recovered by the *Aware Consumer*, which, however, only succeeds because of a fixed retry interval, which we set to 15 s for comparability (notable at the staircase progression). Hence, this scenario is highly susceptible to varying delays and requires round trip time estimation at the consumer or polling, both add overhead. *Lifetime Adoption* on forwarders prevents data loss due to expired PIT state (unlike *Baseline*). Active PIT state enables data forwarding once it is ready, without polling. This mechanism introduces two drawbacks, though: High-speed forwarders are unlikely to adopt extremely long and non-standard timeout values; Interest aggregation prevents the recovery of actual losses while PIT state is active, leading to 20% of loss in our experiments. Overall, end-to-end data retrieval challenges Vanilla ICN scenarios and exhibits high data losses when producer delays are untypically long. Built-in ICN mechanisms can mitigate this effect, if configured properly, but impose strong assumptions on "regular" ICN nodes that are not aware of the delay domain. We overcome these drawbacks using our two ICN extensions.

ICN Extensions: With *Delay-tolerant retrieval*, the majority of requests finishes in almost exactly 32 s, the returned WAIT time of the gateway after the first Interest request. Future work can utilize gateway knowledge more effectively and return a more accurate time value to reduce completion time. This extension enables arbitrary producer delays without the need for long-lived forwarding state. Furthermore, the removal of PIT entries on the first data response of the gateway terminates futile retransmission, compared with Vanilla ICN, and obsoletes round trip time estimation on the Internet application. The producer-initiated *Reflexive push* reflects sporadic IoT data generation most naturally. All transactions finish reliably with 100% success, and the completion times directly reflect the delay distribution of DSME. An additional round trip of the nested double Interest/data flow has a negligible performance overhead when directed toward the fast network. Similarly to *Delay-tolerant retrieval*, this approach enables arbitrary producer delays, and avoids the need to adopt long `InterestLifetime` values on Internet forwarders. In contrast, the unsolicited data push on the wireless link reduces the energy consumption on constrained nodes by a factor of approximately 1.7 when compared to a regular request-response ICN flow.

Unresolved Challenges: LoRa-ICN has contributed vital network mechanics but challenges remain in three key areas. First, there is a need to devise efficient slot negotiation and scheduling for reducing network join time and energy consumption. Second, enabling security is crucial. This encompasses the bootstrapping of LoRa nodes and gateways, a trustworthy registration process, and authenticating LoRa nodes before any gateway can act on their behalf. Third, interfacing of Internet-based applications with distributed LoRa-ICN nodes across sub-networks requires data synchronization techniques and advanced ICN mechanisms, such as publish-subscribe.

CONCLUSIONS AND FUTURE WORK

LoRa is an attractive radio technology for the IoT, providing a long wireless transmission range for battery-driven devices. Its versatility is hindered, though, by common deployments with LoRaWAN.

We re-visited LoRa in the IoT to provide a serverless, data-oriented communication service. We presented the design of a new media access and network layer that leverages 802.15.4 DSME and Information-centric Networking to allow for reliable LoRa transmissions. To scale to a global Internet (of Things), LoRa-ICN facilitates ubiquitous connectivity of constrained nodes and robust bi-directional communication in the presence of power-saving regimes and high loss rates.

We showed that vastly higher latencies in low-power wireless domains can be addressed by extending the default ICN node behavior at the network edge. Two protocol extensions enable ICN-style data transport between resource-constrained LoRa nodes and a domain-agnostic application on the ICN Internet. The core idea is not limited to LoRa but caters to various delay-prone scenarios. Our experiments based on common IoT hardware and software showed significant performance improvements and further optimization potential compared to Vanilla ICN.

The new LoRa-ICN system paves the way for more versatile LoRa deployments in the IoT that serve additional use cases, for example, mixed sensor-actor topologies, or firmware updates utilizing beacon overloading. Future work shall extend our system design by two components, which will improve the performance and reliability. First, advancing the PHY mapping of DSME-LoRa shows promise for exploiting the chirp spread spectrum modulation better, to increase the transmission range or the total throughput. This may involve partial frequency overlapping techniques because LoRa signals can be demodulated below a certain noise level. Second, by exploiting the link-knowledge of the LoRa-ICN gateway, an estimator model can be derived that achieves more accurate timing behavior in our *Delay-tolerant retrieval* extension.

ACKNOWLEDGMENTS

This work was supported in parts by the German Federal Ministry for Education and Research (BMBF) within the projects PIVOT and C-ray4edge, and by the Guangdong Provincial Key Laboratory of Integrated Communications, Sensing and Computation for Ubiquitous Internet of Things and by the Guangzhou Municipal Science and Technology Project 2023A03J0011.

REFERENCES

[1] D. Zorbas *et al.*, "TS-LoRa: Time-Slotted LoRaWAN for the Industrial Internet of Things," *Computer Commun.*, vol. 153, 2020, pp. 1–10.

[2] L. Leonardi, L. Lo Bello, and G. Patti, "MRT-LoRa: A Multi-Hop Realtime Communication Protocol for Industrial IoT Applications Over LoRa Networks," *Computer Commun.*, vol. 199, 2023, pp. 72–86.

[3] IEEE 802.15 Working Group, "IEEE Standard for Low-Rate Wireless Networks," IEEE Std 802.15.4™–2015 (Revision of IEEE Std 802.15.4- 2011), 2016.

[4] L. Seitz *et al.*, "Use Cases for Authentication and Authorization in Constrained Environments," IETF, RFC 7744, Jan. 2016.

[5] E. Baccelli *et al.*, "RIOT: An Open Source Operating System for Low-End Embedded Devices in the IoT," *IEEE Internet of Things J.*, vol. 5, no. 6, Dec. 2018, pp. 4428–40.

[6] J. Finnegan, S. Brown, and R. Farrell, "Evaluating the Scalability of LoRaWAN Gateways for Class B Communication in ns-3," *Proc. IEEE Conf. Standards for Commun. Networking*, IEEE, 2018, pp. 1–6.

[7] Y. Shiferaw, A. Arora, and F. Kuipers, "LoRaWAN Class B Multicast Scalability," *Proc. IFIP Networking Conf.*, IEEE Press, June 2020, pp. 609–13.

[8] L. Zhang *et al.*, "Named Data Networking," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, 2014, pp. 66–73.

[9] P. Kietzmann *et al.*, "Delay-Tolerant ICN and Its Application to LoRa," *ACM Conf. Information-Centric Networking*, ACM, Sept. 2022, pp. 125–36.

[10] J. Alamos, T. C. Schmidt, and M. Wählisch, "6LoRa: Full Stack IPv6 Networking with DSME-LoRa on Low Power IoT Nodes," *Proc. EWSN'23*, New York, USA: ACM, Sept. 2023.

[11] J. Alamos *et al.*, "DSMELoRa: Seamless Long Range Communication Between Arbitrary Nodes in the Constrained IoT," *Trans. Sensor Networks*, vol. 18, no. 4, Nov. 2022, pp. 1–43.

[12] Y. Li *et al.*, "NDN Automatic Prefix Propagation," NDN, Technical Report NDN-0045, Mar. 2018.

[13] M. Król *et al.*, "RICE: Remote Method Invocation in ICN," *Proc. ACM Conf. Information-Centric Networking*, New York, NY, USA: ACM, 2018, pp. 1–11.

BIOGRAPHIES

PETER KIETZMANN is a Ph.D. student of the Internet Technologies research group at the Hamburg University of Applied Sciences. His particular interests lie in radio technologies, embedded programming, as well as protocol and system security for the IoT. Recently, Peter has focused on Information-centric Networking in the context of low-power radio networks of the ultra-constrained IoT.

JOSÉ ALAMOS is a Ph.D. student in Computer Science and a research assistant at the Hamburg University of Applied Sciences. As part of his research, he is designing reliable and low-power LPWAN MAC protocols with a focus on standardization. José is actively developing the RIOT operating system and serves as a maintainer since years.

DIRK KUTSCHER is a professor at the Hong Kong University of Science and Technology in Guangzhou and the co-director of the Future Networked Systems Laboratory at HKUST(GZ). His main interests lie in the intersection of distributed computing and networking for Compute-First Networking and in Internet architecture. Previously Dirk was a professor at the University of Applied Sciences in Emden/Leer, CTO for Virtual Networking at Huawei Research Germany, Chief Networking Researcher at NEC Europe, and Visiting Researcher at KDDI R&D in Japan. He is co-chairing two IRTF Research Groups, ICNRG and DINRG.

THOMAS C. SCHMIDT is a professor of computer networks at the Hamburg University of Applied Sciences, where he heads the Internet Technologies Research Group. He led numerous research projects, was a Visiting Professor with the University of Reading, U.K., and is a co-founder and coordinator of the open source community developing RIOT. His current research interests include development, measurement, and analysis of large-scale distributed systems like the Internet. He served as a co-editor and technical expert on several occasions and is actively involved in the work of IETF/IRTF since many years.

MATTHIAS WÄHLISCH holds the Chair of Distributed and Networked Systems at the Faculty of Computer Science at TU Dresden. In addition to his academic activities, he is a co-founder and a coordinator of some successful open source projects such as RIOT and active in the IETF and IRTF since 2005. He is the recipient of the Young Talents Award of Leibniz-Kolleg Potsdam for outstanding achievements in advancing the Internet, and several best paper awards, for example from ACM CoNEXT 2022.